HHS Public Access

Author manuscript

Neurobiol Learn Mem. Author manuscript; available in PMC 2024 October 28.

Published in final edited form as:

Neurobiol Learn Mem. 2022 November ; 195: 107688. doi:10.1016/j.nlm.2022.107688.

Generalizing the Control Architecture of the Lateral Prefrontal Cortex

McKinney Pitts,

Derek Evan Nee

Department of Psychology, Florida State University, Tallahassee, FL 32306-4301

Abstract

Cognitive control guides non-habitual, goal directed behaviors allowing us to flexibly adapt to ongoing demands. Previous work has suggested that multiple cognitive control processes exist that can be classed according to their action on present-oriented/external information versus future-oriented/internal information. These processes can be mapped onto the lateral prefrontal cortex (LPFC) such that increasingly rostral areas are involved in increasingly future-oriented/internal control processes. Whether and how such processes are organized to support goal-directed behavior remains unclear. On the one hand, the LPFC may flexibly adapt based upon demands. On the other hand, there may be a consistent control architecture such as a control hierarchy that generalizes across demands. Previous work using fMRI in humans during a comprehensive control task that engaged several control processes at once found that an area in mid-LPFC consistently exerted widespread influence throughout the LPFC. These data suggested that the mid-LPFC forms an apex of a putative control hierarchy. However, whether such an architecture generalizes across tasks remains to be tested. Here, we utilized a modified comprehensive control task designed to alter how control processes influence one another to test the generalizability of the LPFC control architecture. Univariate fMRI activations revealed distinct control-related activations relative to past work. Despite these changes, effective connectivity modeling revealed a directed architecture similar to previous findings with the mid-LPFC exerting the most widespread influences throughout LPFC. These results suggest that the fundamental control architecture of the LPFC is relatively fixed, and that different demands are accommodated through modulations of this fixed architecture.

Keywords

cognitive control; hierarchy; prefrontal cortex; fMRI; dynamic causal modeling

INTRODUCTION

Cognitive control is the ability to organize thought and behavior in accordance with internal goals (Miller and Cohen 2001). Cognitive control is utilized in our daily lives in many ways such as: scheduling multiple concurrent tasks, proactively preparing for the future, or breaking habitual responding in favor of more intentional behavior. These processes

Corresponding Author: Derek Evan Nee, Florida State University, 1107 W Call St, Tallahassee, FL 32306-4301, nee@psy.fsu.edu.

organize behavior to deal with complex problems as they arise. Understanding these fundamental functions is necessary to better elucidate how to optimize control and remediate control when it goes awry.

Cognitive control engages the frontoparietal network (FPN) (Cole and Schneider, 2007; Dosenbach et al., 2007; Vincent et al., 2008; Cole et al., 2013). Therein, the lateral prefrontal cortex (LPFC) is known to be essential in performing many higher-level cognitive abilities such as planning, attention, and working memory (Bechara et al., 1994; Miller & Cohen, 2001; Tanji & Hoshi, 2001; Tanji & Hoshi, 2008). However, the potential organization of specific control processes within these regions and how they support higher-level cognition remains uncertain.

According to the Adaptive Coding Model, the FPN along with other areas are engaged in diverse situations to organize cognitive processes towards intelligent cognition (Duncan, 2001; Duncan, 2010; Duncan et al., 2020). Under this framework, control processes do not have a systematic organization, but rather, flexibly adapt to current demands (Duncan, 2013). By contrast, several studies have provided evidence that the processes underlying cognitive control are systematically organized (Koechlin et al., 2003; Nee & Brown, 2012; Bahlmann et al., 2014; Nee & D'Esposito, 2016), although specifying a proper organizing framework has been elusive (Koechlin & Summerfield, 2007; Badre, 2008; Badre & Nee, 2018). Recently, we suggested that the FPN, and LPFC within it, are broadly organized such that somatomotor proximal areas act upon present-oriented, external representations while somatomotor distal areas act upon future-oriented, internal representations (Nee, 2021). Within this framework, we concern ourselves here with three different putative control processes: In the immediacy, learned stimulus-response pairings inform correct action selection. Selecting amongst them is referred to as *sensorimotor control*, and it governs the most concrete contributions to behavior that differentiate habitual responding from intentional control. These mappings can in turn be guided by current task sets and contextual rules. *Contextual control* incorporates the context to select and maintain task sets that inform proper stimulus-action selection. Finally, context-task set mappings may themselves flexibly change over different temporal epochs. *Episodic control* is the contribution of control associated sustaining a temporal episode to inform context-task set mappings (Koechlin et al., 2003; Koechlin and Summerfield, 2007). These processes are thus organized by both timescale (Fuster, 2001; Koechlin et al., 2003; Koechlin & Summerfield, 2007) and abstraction (Badre, 2008).

By definition, control denotes source-target relationships such that a systematic organization of control processes begets a systematic ordering of influences. For example, in the framework described above, episodic control can influence contextual control which, in turn, can influence sensorimotor control. This chain of influences has been formalized under the Cascade Model which posits a hierarchical ordering of LPFC areas and control processes (Koechlin et al., 2003; Koechlin & Summerfield, 2007). By contrast, a heterarchical organization with no single control process/brain area being generally dominant over others would be predicted by a framework that posits that control processes are not systematically organized, such as the Adaptive Coding Model (Duncan, 2010; Crittenden and Duncan, 2014). Note that the Adaptive Coding Model does not rule out that the LPFC or control

more generally acts hierarchically under certain demands, but rather, that distinct demands would entail distinct organizations such that no one organization would generally persist.

Pertinent to this debate, under accounts that subscribe to a putative hierarchical organization of the LPFC, the organization itself remains a point of contention. The Cascade Model predicts that the rostral-most part of the LPFC, the lateral frontal pole (FPl), sits atop the control hierarchy as the apex (Koechlin & Summerfield, 2007). Badre & D'Esposito (2009) reviewed evidence describing the anatomical plausibility of an FPl apex. These authors formalized that hierarchies can be operationally defined through dominance relationships such that an apex would be marked by asymmetrically greater influences upon other areas of the brain than it receives. Using this formalization, Nee & D'Esposito (2016) modeled effective connectivity (directed influences) in the LPFC during a comprehensive control task that factorially manipulated multiple forms of control. In contrast to the predictions of an FPl apex, they found that an area in the mid-LPFC showed the maximal outward asymmetries of influence that are consistent with an apex. Remarkably, the FPl was the least influential of the LPFC areas modeled. These observations were further bolstered by causal data using transcranial magnetic stimulation (Nee & D'Esposito, 2017), and are broadly consistent with anatomical projections in the monkey (Goulas et al., 2014). However, while the task of Nee & D'Esposito manipulated sensorimotor and contextual control, at the highest level, the task engaged what they referred to as *temporal control*. This was defined as the maintenance of internal representations to guide future cognition. Although temporal control and episodic control both concern themselves with the common requirement of sustaining representations for the future, only in episodic control do those representations act directly upon task sets (i.e. contextual control). On the other hand, temporal control requires no direct interaction among internalized representations and ongoing task sets. Hence, the lack of influence of the FPl upon the LPFC in the study by Nee & D'Esposito may have resulted from a failure to engage the proper control process. Moreover, if the engagement of episodic control causes FPl to act in an apical manner, this suggests that the control hierarchy flexibly shifts according to demands. Such data might be consistent with the idea that there is no one hierarchical control architecture, but rather, an adaptive control architecture that appears to be heterarchical when examined across diverse demands.

Here, we sought to investigate whether a single hierarchical control architecture generalizes across different tasks or whether control architectures flexibly shift according to demands. We implemented a modified version of the comprehensive control task employed by Nee & D'Esposito (2016) by replacing temporal control with episodic control. Critically, this manipulation shifts the influence of control processes among one another providing an opportunity to alter the hierarchical organization of cognitive control. We hypothesized that the inclusion of episodic control may allow the FPl to act in an apical manner. Moreover, such a shift would also provide evidence for the Adaptive Coding Model by demonstrating that no one hierarchical organization of the LPFC exists, but rather, the organization adapts to diverse demands. However, if the mid-LPFC remains apical, it would suggest that the LPFC has a general hierarchical control architecture. Moreover, such a result would suggest that the mid-LPFC, as a general apex of cognitive control, would make for an important target for interventions seeking to broadly modulate control processing in diverse situations.

METHODS

Participants

Data from thirty subjects were collected. Technical issues corrupted the data of 4 participants, and 1 participant was excluded for excessive movement (>0.25 mm root mean squared frame-wise displacement (RMS FD); mean RMS FD of remaining sample = 0.1062 mm). The analyzed sample included twenty-five right-handed native English speaking Florida State University students (average age= 19.48, 18 female). The sample size was based upon past studies (Nee & D'Esposito, 2016; 2017) which showed power sufficiency evidenced by empirically validated replicability at the chosen sample size (Nee, 2019). All reported participants completed an initial screening which included demonstrating the ability to understand and perform the task to satisfactory levels. Informed consent was obtained in concordance with the Office for Human Subjects Protection (OHSP) & FSU Institutional Review Board (IRB).

Procedure and Paradigm

The paradigm was written in MATLAB R2021a (Mathworks, Inc.) with Psychophysics Toolbox Version 3 (Brainard, 1997; Pelli, 1997; Kleiner et al, 2007). The task was a modified version of the paradigm used by Nee & D'Esposito (2016; 2017) designed to alter the interdependency among control processes. Each trial consisted of the presentation of a letter in one of five locations on a 5-pointed star (Fig 1A). Letters were drawn from the word "T-A-B-L-E-T". Each letter was surrounded by a colored frame of a given shape. For all trials in each block, either the letters (verbal trials) or spatial locations (spatial trials) were relevant, indicated by the color of the frame around the letter (Fig 1B). Participants responded to each stimulus with a yes/no keypress. The correct response depended upon the task being performed. On sequence-start trials, participants responded regarding whether the stimulus began the sequence ("T" for verbal trials; top of the star for spatial trials). On sequence-back trials, participants responded regarding whether the current stimulus followed the previous stimulus in the sequence (e.g., "B" following "A" for the verbal trials; left point following lower right point on spatial trials). Diamond-shaped frames cued the sequence-start task, while square-shaped frames cued the sequence-back task (Fig 1C).

Sequence-start and sequence-back formed the task sets in the current study. We define a task set as a finite mapping of input-output associations, such that outputs can be specified abstractly (e.g. "yes"/"no") rather than concretely (e.g. left index/right index). Similarly, inputs can be specified categorically (e.g. start/non-start) rather than concretely tied to a particular stimulus (e.g. "T", top left, etc.). By this definition, the same task sets can be flexibly deployed to both verbal and spatial trials.

All blocks began with a sequence-start trial followed by 1-2 sequence-back trials (baseline 1). Baseline blocks continued with an additional 6-8 sequence-back trials. These blocks placed consistent demands on *sensorimotor control* to select the appropriate response based upon the stimulus. Other blocks were split into three phases: baseline 1, subtask, and baseline 2, wherein the subtask phase added additional control demands. The transition from

baseline 1 to subtask was indicated by a subtask cue (Fig 1D,E), alerting participants to the forthcoming demands:

On switch blocks, a subtask cue (triangle frames) indicated to the participant that they would need to randomly switch between sequence-start (diamond frames) and sequence-back (square frames) tasks. Compared to baseline blocks, switch blocks required *contextual control* to use the context (frame shape) to select the appropriate task set (sequence-start vs sequence-back).

On delay blocks, the subtask cue indicated a task set mapping that was specific to the block (e.g. circle -> sequence-back, plus -> sequence-start or the converse). These blocks engaged *episodic control* to maintain context-task set mappings specific to the block (i.e. episode). To minimize ongoing contextual control demands, the subtask phase on delay blocks consisted solely of sequence-back trials cued by square frames. Only a single critical trial required the use of the task set mapping (return trial), which was followed by a return to sequence-back trials (baseline 2).

Finally, dual blocks combined the requirements of the switch and delay conditions. As in switch blocks, participants switched between sequence-start and sequence-back tasks thereby engaging contextual control. As in the delay blocks, the subtask cue indicated an episode-specific task set mapping thereby engaging episodic control. Critically, only on dual blocks, all trials in the subtask phase included either circle- or plus-shaped frames. This required participants to use the sustained task set mapping to select and execute the appropriate task throughout the subtask phase. In this way, dual blocks required *contextual control informed by episodic control*.

Switch, delay, and dual blocks ended with 1-2 trials of sequence-back (baseline 2). Focusing on the subtask phases, baseline blocks contained low contextual control and low episodic control, delay blocks contained low contextual control and high episodic control, switch blocks contained high contextual control and low episodic control, and dual blocks contained high contextual control and high episodic control.

Collectively, the paradigm was designed as a 2x2x2 factorial with factors of stimulus domain (spatial, verbal), contextual control (low, high), and episodic control (low, high). Each stimulus was presented for 500 ms separated by a uniformly jittered interstimulus intervals of 2600-3400 ms in steps of 400 ms. At the end of each block, feedback that showed the number of correct trials compared to the total trials of the block was presented for 500 ms followed by an interblock interval of 2600 or 3400 ms. After a run was finished, a percentage accuracy was given for the current run, for the previous run, and for the total amount of runs completed thus far. A short break was administered before the next run began. Participants were asked to always maintain fixation on a cross within the center of the screen and were visually monitored using an EyeLink 1000 eye tracker (SR Research Ltd.) for adherence to these instructions, and for vigilance. The experimental session consisted of 6 runs of 144 trials each for a total of 864 trials per subject. The mappings of frame colors, frame shapes, and response keys were all counterbalanced at the participant level. Given the complexity of the paradigm, experimental sessions were preceded by a learning and practice

session (average time between sessions= 5.12 days). Only participants who demonstrated the ability to understand the instructions and perform the task completed the experimental session.

Many of the details of the present paradigm are comparable to the paradigm of Nee and D'Esposito (2016; 2017). The most substantial differences owe to the control demands in the delay and dual conditions. Previously, the delay and dual conditions required participants to retain an item (e.g. a letter for verbal trials) throughout the subtask phase which then informed a sequence-back decision on the return trial (e.g. does the present letter follow the retained letter in the sequence). By contrast, the delay and dual conditions of the present paradigm require the retention of a cue-task set mapping which informs task set selection on the return trial. In this way, our previous paradigm placed demands on what we referred to as "temporal control", while the present paradigm places demands on episodic control. In both cases, the control demands operate over epochs of several seconds such that each demand is more temporally abstract than control demands that operate over shorter intervals. Hence, if control demands are purely a function of time (i.e. temporal abstraction), one might expect episodic and temporal control to engage the same neural systems. If, by contrast, control varies as a function of the representation being acted upon (i.e. selecting a task set versus determining sequential order), one would expect different neural systems to be engaged by each form of control. Notably, the term "temporal control" was used previously as a term agnostic to whether or not the form of control would be different from other forms of control that operate across sustained temporal epochs (Nee and D'Esposito, 2016).

Image Acquisition

Images were acquired on a Siemens 3T Prisma with a 32-channel head coil located in the Magnetic Resonance Imaging facility within Florida State University's College of Medicine. The participant viewed stimuli while in the scanner by way of a head coil attached mirror reflecting onto a projection screen at the bore of the magnet. Responses were collected by way of a 4-button MR compatible button box (Current Designs, Inc.) using the index finger of each hand.

Functional T2*-weighted images were acquired using an EPI sequence with 35 descending slides and 3.4x3.4x3.8mm voxels (TR= 2000ms; echo time= 25ms; flip angle= 70; field of view= 220). Four dummy scans preceded each functional scan to allow for image stabilization. Phase and magnitude images were collected to estimate the magnetic inhomogeneity. High- resolution T1-weighted MPRAGE images were collected for spatial normalization (384x384x256 matrix of 0.667mm^3 isotropic voxels; TR= 1840ms; echo time= 2.9ms; flip angle= 9).

Image Preprocessing

Preprocessing was performed using SPM12 (http://www.fil.ion.ucl.ac.uk/spm/) unless otherwise specified. Raw data were converted from DICOM to NIfTI format. Origins for all images were manually set to the anterior commissure. Functional data were then spike-corrected to reduce the impact of artifacts and outliers using AFNI's 3dDespike routine (http://afni.nimh.nih.gov/afni). Images were corrected for differences in slice timing

on a per-run basis using sinc-interpolation. Head motion was corrected using a rigid body transformation. Images were unwarped by utilizing the FieldMap toolbox, which also corrected motion-by-susceptibility distortions (Andersson et al.,2001). Functional data were then coregistered to the T1 structural scan from which tissue probability maps were calculated (Ashburner and Friston, 1997). The tissue probability maps were then used to warp the individual to the MNI152 brain template. As part of the spatial normalization the data were resampled to 2 mm^3 isotropic. 4-mm full-width/half-maximum isotropic Gaussian smoothing was applied to all spatially normalized functional images.

Univariate Image Analysis

All analyses were performed using SPM12 and MATLAB R2021a. Each participant's data was fit with a general linear model that included separate regressors for each control condition (baseline, switch, delay, and dual) crossed by each stimulus type (spatial, verbal). These eight regressors spanned the subtask phase forming the regressors of main interest. Since baseline blocks did not have a separate subtask phase, a middle phase of the block was modeled that was matched to the other blocks in durations. Transient regressors of non-interest separately modeled the initial trial in each block, the subtask cue that signaled each subtask, and the return trial. Sustained regressors of non-interest modeled the baseline 1 and baseline 2 phases. Additional transient regressors for left and right button presses crossed with task (sequence-start, sequence back), and errors were also included to remove variance of non-interest from the sustained regressors. Each regressor was convolved with a canonical hemodynamic response function. 14 motion regressors (6 linear, 6 squared linear, 1 framewise displacement, 1 squared framewise displacement) were included to account for artifacts caused by motion. Moreover, any frame with 0.5 mm or greater framewise displacement was censored. A high-pass filter of 192s was included to remove low frequency trends, and AR(1) modeling was included to correct for temporal autocorrelation.

Separate contrasts were performed at the first level for the main effect of contextual control (dual+switch > delay+baseline) and the main of episodic control (dual+delay > switch+baseline). These contrasts were then submitted to respective second-level one sample t-tests. Whole-brain maps were thresholded at $p < 0.001$ at the voxel level with cluster extent (89 for contextual, 115 for episodic control) providing family-wise error correction according to Gaussian random field theory as implemented in SPM12. Since the algorithm depends upon the smoothness of the data (residuals) which varied as a function of second-level models, different extents were calculated for different effects. As estimated by SPM12, the smoothness of the episodic control data was 8.4 x 8.6 x 8.5 mm and the smoothness of the contextual control data was 8.1 x 8.3 x 8.1 mm.

Dynamic Causal Modeling

In order to investigate directional influences (effective connectivity) among LPFC areas, we turned to dynamic causal modeling (DCM) (Friston et al. 2003) implemented in DCM12 within SPM12. DCM provides a biophysically plausible generative model of the time series of regions-of-interest (ROIs). Here, we performed bilinear DCM that models each region's time series as a function of A) fixed effective connectivity among brain regions,

B) modulations of effective connectivity as a function of experimental manipulations, and C) exogenous perturbations (stimulus inputs). Interactions among brain regions is modeled at the neuronal level with a hemodynamic model specifying the transfer function from neuronal to hemodynamic activity (see Stephan et al., 2010; Zeidman et al., 2019 for accessible reviews).

Following Nee & D'Esposito (2016; 2017), our model consisted of 6 regions spanning the rostral-caudal and dorsal-ventral axes of the LPFC: the lateral frontal pole (FPl), middle frontal gyrus (MFG), inferior frontal sulcus (IFS), caudal middle frontal gyrus (cMFG), superior frontal sulcus (SFS), and inferior frontal junction (IFJ) (Fig 2A). All ROIs were selected from the left hemisphere. This decision was motivated by meta-analyses of the hierarchical control literature which has demonstrated a strong left lateralization for the control processes under study (Nee et al., 2014; Badre and Nee, 2018), and that the verbal-ventral/spatial-dorsal distinction is present only in the left hemisphere. Given the variability of LPFC togography (Braga and Buckner, 2017) and corresponding correlates of cognitive control (Smith et al., 2021), ROIs were localized within each individual. Following our past work, caudal ROIs were defined using stimulus domain contrasts: the spatial > verbal contrast for SFS and the verbal > spatial contrast for the IFJ. Middle ROIs (cMFG and IFS) were isolated by utilizing the contextual control contrast (dual+switch > delay+baseline). Rostral ROIs (FPl and MFG) were identified using the dual > baseline contrast. The dual > baseline contrast was chosen over the episodic control contrast (dual+delay > switch+baseline) given that the delay condition did not uniquely activate any LPFC area more than other control conditions thereby diluting the main effect contrast. An initial search for each peak began near the previous peaks of Nee & D'Esposito (2016) (FPl: [−44,48,4]; MFG: [−38,28,44]; IFS: [−52,20,28]; cMFG: [−34,10,60]; SFS: [−22,0,54]; IFJ: [−38,6,26]). Peaks were defined as local maxima as indicated by SPM12 in the vicinity of the previously reported group means and visually inspected to adhere to gyral/sulcal anatomy. Peaks were originally searched for at a threshold of $p<0.001$ (69.33% of clusters), and if no cluster was found, the threshold was progressively lowered until a suitable peak was identified (0.01 – 14.67% of clusters, 0.05 – 7.33% of clusters, 0.1 – 1.33% of clusters, 0.5 – 7.33% of clusters). If no peak could be identified, the ROI was centered at the group average (4.67% of clusters). Each region was then defined as a 6 mm radius sphere centered upon each activation peak. Suprathreshold data from each sphere was extracted and corrected for confounds in the same manner as the univariate analyses (whitened, high-pass filtered, motion regressed, and high movement frames censored). As previously reported, the different contrasts are not equally powered (Nee, 2019), which can result in differences in the number of voxels for different ROIs after thresholding. Accordingly, we observed that some ROIs had more voxels than others $F(1,24)= 2.88$, $p= 0.0173$). However, this effect was driven by contrast definition as no pair of areas defined by the same contrast had different numbers of voxels (all $p > 0.2$). A corollary is that there was no difference between our critical areas of interest: MFG and FPl ($t(24) = −1.27$, $p = 0.22$). Moreover, these ROIs are ultimately reduced to summaries as described below.

Each regional time series was summarized using the first eigenvariate of voxels comprising the region thereby forming the ROI time series for DCM. To ensure that these summaries properly reflected the signals in the ROI, amount of voxel-wise variance explained by the

first eigenvariate was examined. These were generally high (mean for each area > 70%; range 70.6-78.6%), but did vary as a function of region ($F(1,24)= 3.34$, $p= 0.0074$). This variability was related to the size of ROIs described above as there was a tendency for the first eigenvariate to explain less variance for larger ROIs ($r = −0.15$, $p = 0.05$ after residualizing out subject effects). Critically, there was no difference between MFG and FPl ($t(24)= −1.09$, $p= 0.29$). Ultimately, the ROI selection procedures, which were designed to tailor ROIs to the individual and capture as much signal of interest while removing noise, did not result in observable signal differences in the MFG and FPl that could influence their ability to drive signals in other areas.

DCM is a model comparison technique. We compared ten models each with different assumptions of how experimental manipulations modulated effective connectivity, and the locus of exogenous perturbations. The fixed connectivity architecture was kept constant across models and was based upon considerations of anatomical connections and functional connectivity as detailed previously (Nee & D'Esposito, 2016;2017). All models also assumed that letter/spatial inputs entered the system at the caudal-most areas, consistent with the dense connections of these areas to posterior areas involved in stimulus representation (Petrides, 2005). Spatial and verbal inputs were chosen to enter via the dorsal SFS and ventral IFJ, respectively due to the well-documented dorsal-ventral split between spatial and verbal information (Goldman-Rakic, 1987, Romanski, 2004) and validated via the stimulus domain contrasts. From this common architecture, we compared five different models of modulations of effective connectivity:

1. A Null Model wherein experimental manipulations did not alter effective connectivity (Fig 2B)
2. A Cascade Model wherein episodic control modulated effective connectivity from rostral-to-mid areas, and contextual control modulated effective connectivity from mid-to-caudal areas (Koechlin et al., 2003; Koechlin and Summerfield, 2007). In other words, cognitive control related modulations of effective connectivity cause a rostral-to-caudal cascade of influences upon action. (Fig 2C)
3. A Modified Cascade Model that added caudal-to-mid stimulus domain modulations to 2). Although the original Cascade Model is a domain-general model of cognitive control, it had not been applied across distinct stimulus domains. Given prior work indicating that models that include stimulus domain modulations were preferred (Nee & D'Esposito, 2016), we examined whether simply adding stimulus domain modulations to the Cascade Model would provide a superior model fit. (Fig 2D)
4. The model described by Nee & D'Esposito (2016) (N&D2016). In this model, rostral-to-caudal cognitive control-related influences and caudal-to-rostral stimulus domain-related influences converge in middle areas. (Fig 2E)
5. The model described by Nee & D'Esposito (2017) (N&D2017). A variant of 4) which adds bidirectional modulations of connectivity among MFG and SFS. (Fig 2F)

In addition to the five models of modulations, we varied how cue inputs perturbed the system. Previously, it was assumed that cues entered the system through the MFG consistent with its role in representing task rules (Nee & Brown, 2012; Bunge et al., 2005; Genovesio et al., 2005; Wallis et al., 2001). This choice was bolstered by examinations of univariate activations and preliminary DCM (Nee & D'Esposito, 2016). However, based on theoretical ideas that the FPl may provide a relay for internally generated schematic representations (Badre & Nee, 2018), which might include higher-order task rules (Nee & Brown, 2012), we formally tested whether models wherein cues entered via the FPl provided superior fits. Hence, the five models described above were crossed with two models of cue inputs (cue -> MFG, cue -> FPl) (Fig 2F) creating a space of ten models.

Random-effects Bayesian model selection (BMS) was used to adjudicate among proposed models (Stephan et al., 2009). First, the models were grouped into two families as a function of the presumed locus of cue inputs, and the model families were compared (Penny et al., 2010). Next, the five models within the winning family were directly compared, and also compared as three families (null, Cascade, N&D). A parametric empirical bayes (PEB) model was constructed from the pooled parameters of each subject's runs for the winning model (Zeidman et al., 2019). PEB allowed us to account for the variance of parameters between runs and subjects and weight connections according to their reliability. These PEB models were then carried to the next level where an additional PEB model upon the subject level models was performed to create a group level model. Inference on connections was performed via one-sample t-tests. Connections with a p-value less than 0.05 after stepwise false-discovery rate (FDR) correction were considered to be significant.

To formally compare the present results with those observed previously, the DCM PEB procedures were repeated on the Nee & D'Esposito (2016; 2017) datasets, using processing reported in Nee (2021). The major difference from prior reports is that a reduced smoothing kernel (4 mm) was included as part of preprocessing, matching the present preprocessing pipeline. Differences in connections comparing the present dataset with Nee & D'Esposito (2016), and Nee & D'Esposito (2017) were examined using the PEB framework. Differences that replicate across studies are reported.

Hierarchical Strength

To examine asymmetries of influence that are indicative of hierarchical rank (Badre & D'Esposito, 2009) we calculated hierarchical strength. Hierarchical strength was defined as the magnitude of participant's efferent connections to a region minus the magnitude of their afferent connections. To determine magnitude, each subject-level connection was multiplied by the sign of the group mean. Connections within the same level of the rostro-caudal hierarchy were excluded (e.g. SFS to IFJ) as these connections were thought to be hierarchically equivalent with the potential to dilute rostral-caudal patterns (Nee & D'Esposito, 2016; 2017). An area that is more apical than others would show the highest positive value on this metric. Notably, these procedures differ slightly from our past work (Nee & D'Esposito, 2016) which rectified effective connectivity. However, stronger magnitude negative influences suggested that negative signs should be taken seriously, hence

the change in procedures which matches procedures we have used more recently (Nee, 2021).

RESULTS

Behavioral Results

We began by verifying that the paradigm produces dissociable behavioral signatures of cognitive control. 2x2x2 ANOVAs with factors of episodic control, contextual control, and stimulus domain were performed on error rates (ER) and reaction times (RT) of correct trials (Fig 3). We focus on subtask trials wherein the cognitive control demands were manipulated.

These analysis revealed significant main effects of episodic control (ER: $F(1,24)= 34.01$; $p<0.001$; RT: $F(1,24)= 31.24$, $p<0.001$) and contextual control (ER: $F(1,24)= 65.4$; $p<0.001$; RT: $F(1,24)= 46.07$, $p<0.001$). The interaction between episodic and contextual control was significant in ER ($F(1,24)= 5.62$; $p=0.0261$), but not RT ($F(1,24)= 0.14$, $p=0.7113$). A significant effect of stimulus domain was found in both ER and RT (ER: $F(1,24)= 9.41$; $p=0.0053$; RT: $F(1,24)= 9.5$, $p=0.0051$) with overall poorer performance for verbal trials. However, stimulus domain did not significantly interact with either form of control (ER: contextual * stimulus domain $F(1,24)= 0.08$, $p= 0.7735$, episodic * stimulus domain $F(1,24)= 0.23$, $p= 0.6373$; RT: contextual * stimulus domain $F(1,24)= 0.06$, $p= 0.8128$, episodic * stimulus domain $F(1,24)= 3.18$, $p= 0.0871$).

Additional ANOVAs were performed on error rates, and reaction times of correct return trials (Fig S1). These analyses revealed significant main effects of episodic control in both error rate and response time (ER: $F(1,24)= 62.43$; $p<0.001$; RT: $F(1,24)= 19.66$, $p<0.001$) and contextual control only in reaction time (ER: $F(1,24)= 0.13$; $p=0.7209$; RT: $F(1,24)= 6.75$, $p=0.0158$). The main effects of control were qualified by an interaction between episodic and contextual control which was significant in ER ($F(1,24)= 19.38$; $p<0.001$), and RT ($F(1,24)= 9.16$, $p=0.0058$). In contrast to the subtask trials in which the interaction among control processes was over-additive in ER, the interactions on return trials were under-additive. The most likely cause of this pattern is that the consistent integration of episodic control and contextual control during the subtask phase of the dual condition (i.e. using the subtask cue to select the task set) facilitated that same process on return trials. This facilitation was absent in the delay condition wherein episodic control did not inform contextual control until the return trial itself. A significant effect of stimulus domain was found in RT but not ER (ER: $F(1,24)= 0.75$; $p=0.3951$; RT: $F(1,24)= 10.94$, $p=0.003$) with slower reaction times for verbal trials. Stimulus domain interacted with contextual control in RT, but no other interactions with stimulus domain were observed (ER: contextual * stimulus domain $F(1,24)=6.86$, $p= 0.0151$, episodic * stimulus domain $F(1,24)= 0.01$, $p= 0.906$; RT: contextual * stimulus domain $F(1,24)= 1.19$, $p= 0.2861$, episodic * stimulus domain $F(1,24)= 0.01$, $p= 0.939$). Given the superior power of the subtask trials by virtue of trial yield, we focus on this phase for future analyses as we have done elsewhere (Nee & D'Esposito, 2016; 2017; D'Mello et al., 2020; Nee, 2021).

The behavioral data verify that the experimental paradigm introduces multiple forms of control as designed. To understand their neural underpinnings, we next examined the fMRI data.

Whole-Brain Results

Next, we examined effects of control on brain activations. Previous work has demonstrated that increasingly abstract control demands engage progressively rostral areas of the LPFC (Koechlin et al, 2003; Badre and D'Esposito, 2007; Nee & Brown, 2013; Bahlmann et al., 2014; Nee & D'Esposito, 2016; see Badre and Nee, 2018 for a review). However, this is not always the case (Reynolds et al., 2012; Pischedda et al., 2017), and the situations that give rise to control gradients or lack thereof remain presently unclear. To examine whether the present paradigm produced clearly differentiable neural correlates of control, we examined the orthogonal main effect contrasts of contextual control and episodic control. Contextual control and episodic control engaged widespread, highly overlapping areas of the LPFC with no clear differentiation along the rostral-caudal axis (Fig 4, **top**). Although the contextual control contrast resembled our previous findings in a related paradigm (Nee & D'Esposito, 2016; 2017), the episodic control contrast much more clearly resembled the contextual control contrast than the more abstract temporal control contrast as previously defined (Fig 4, **bottom**).

To quantify these observations, we computed voxel-wise correlations among the unthresholded t-maps using the contextual control and episodic control contrasts of the present study, and the contextual control and temporal control contrasts of Nee & D'Esposito (2016). As expected, the contextual control contrast in the present study was more similar to the contextual control contrast in Nee & D'Esposito ($r = 0.68$) than the temporal control contrast ($r = 0.45$; difference $Z = 232.20$, $p < 0.000001$). This same pattern was observed when comparing the episodic control contrast to the contextual control contrast of Nee & D'Esposito ($r = 0.63$) and the temporal control contrast of Nee & D'Esposito ($r = 0.32$; difference $Z = 274.03$, $p < 0.000001$). Moreover, while the contextual control and episodic control contrasts in the present study were highly similar ($r = 0.81$), the contextual control and temporal control contrasts of Nee & D'Esposito were much more dissimilar ($r = 0.32$; difference $Z = 541.76$, $p < 0.000001$). Notably, the lack of correlation with the temporal control contrast is not due to noisiness in the contrast map – comparing the temporal control contrast of Nee & D'Esposito (2016) to the same contrast of Nee & D'Esposito (2017) revealed high similarity ($r = 0.74$), while the temporal control contrast of Nee & D'Esposito (2016) was dissimilar to the contextual control contrast of Nee & D'Esposito (2017) ($r = 0.38$; difference $Z = 378.65$, $p < 0.000001$).

Collectively, these data indicate that what is captured by episodic control in the present study shares marked similarity to contextual control both here and previously observed. However, these gross observations may overlook subtle differentiations among these processes at a regional level. So, we turned to ROI analyses next.

Regions of Interest Analysis

6 mm radius spherical ROIs were centered around peaks reported by Nee & D'Esposito (2016) providing an unbiased estimate of effect sizes. We examined contrast estimates in six ROIs spanning the rostral-caudal and dorsal-ventral axes: the lateral frontal pole (FPl), middle frontal gyrus (MFG), inferior frontal sulcus (IFS), caudal middle frontal gyrus (cMFG), superior frontal sulcus (SFS), and inferior frontal junction (IFJ). An ANOVA including factors of stimulus type (verbal, spatial), episodic control (high, low), contextual control (high, low), and region of interest (SFS, IFJ, cMFG, IFS, MFG, FPl) revealed significant main effects of all of factors (stimulus type: $F(1,24)= 4.59$, $p=0.0424$; contextual control: $F(1,24)= 94.5$, $p<0.001$; episodic control: $F(1,24)= 21.89$, $p<0.001$; region: $F(1,24)= 2.37$, $p=0.0432$), as well as interactions among control demands (episodic control*contextual control: $F(1,24)= 6.46$, $p=0.0179$). The interaction was driven by an overall tendency for control effects to be under-additive. However, each experimental main effect interacted with region (stimulus domain*region: $F(1,24)= 36.49$, $p<0.001$; episodic control*region: $F(1,24)= 4.52$, $p<0.001$; contextual control*region: $F(1,24)= 6.98$, $p<0.001$) suggesting some heterogeneity of activations throughout the LPFC.

To further investigate this heterogeneity, subsequent ANOVAs were performed within each ROI. All p-values are reported after Bonferroni correction for multiple comparisons. Replicating our previous reports (Nee & D'Esposito, 2016; 2017) middle and caudal ROIs tended to be sensitive to stimulus domain (IFS: $F(1,24)= 17.24$, $p_{corrected}= 0.0021$; SFS: $F(1,24)= 35.88$, $p_{corrected}<0.001$; IFJ: $F(1,24)= 41.96$, $p_{corrected}<0.001$), but rostral areas were not (MFG: $F(1,24)= 1.17$, $p_{corrected}=1.00$ FPl: $F(1,24)= 0.96$, $p_{corrected}=1.00$). Such data are mostly consistent with the idea that middle and caudal areas have domain preferences, while rostral areas act domain-generally. Ventral mid (IFS) and caudal (IFJ) areas were activated more during verbal than spatial trials, while dorsal caudal SFS showed the opposite. Interestingly, whereas we previously observed a spatial preference in dorsal mid cMFG, an opposite trend was observed here, which did not survive correction (cMFG: $F(1,24)= 4.81$, $p_{corrected}= 0.2268$. The reasons for this reversal are unclear. Nevertheless, stimulus domain preferences replicated our previous work in 5 of the 6 ROIs suggesting reasonable comparability.

Turning next to control-related activations, each ROI showed a significant effect of contextual control (SFS: $F(1,24)= 17.75$, $p_{corrected}=0.0018$; IFJ: $F(1,24)= 42.75$, $p_{corrected}<0.001$, cMFG: $F(1,24)= 31.36$, $p_{corrected}<0.001$, IFS: $F(1,24)= 32.65$, $p_{corrected}<0.001$, MFG: $F(1,24)= 17.04$, $p_{corrected}=0.0024$, FPl: $F(1,24)= 56.74$, $p_{corrected}<0.001$). Effects of episodic control were somewhat weaker, but significant in all areas except the MFG and SFS (SFS: $F(1,24)= 8.23$, $p_{corrected}=0.0510$; IFJ: $F(1,24)= 12.43$, $p_{corrected}= 0.0102$, cMFG: $F(1,24)= 13.81$, $p_{corrected}= 0.0066$, IFS: $F(1,24)= 18.41$, $p_{corrected}= 0.0018$, MFG: $F(1,24)= 2.08$, $p_{corrected}=0.97$, FPl: $F(1,24)= 16.14$, $p_{corrected}= 0.0030$). These effects were qualified by under-additive interaction between episodic and contextual control observed in the MFG and SFS (MFG: $F(1,24)= 12.57$, $p_{corrected}= 0.0096$; SFS: $F(1,24)= 11.05$, $p_{corrected}= 0.0168$; Fig 5).

Finally, no area demonstrated an interaction between stimulus domain and control that passed correction suggesting that control and stimulus domain were largely orthogonal factors in the present dataset.

Collectively, the effects of control did not show a clear ordering as a function of spatial topography. Although there was some tendency for control-related activations to increase in increasingly rostral areas, general reductions in MFG deviated from that pattern (Fig 6). Nevertheless, control-related activations tended to peak in FPl, which showed stronger contextual control ($F(1,24)= 58.17$, $p_{corrected}<0.001$) and episodic control ($F(1,24)= 12.05$, $p_{corrected}=0.002$) related activations relative to MFG. The strong activations in FPl may reflect a role in driving widespread LPFC activations, consistent with an apical role (Badre and D'Esposito, 2009; Koechlin and Summerfield, 2007). However, such a role cannot be inferred from activations alone. To further investigate the strength and direction of how these regions communicate, we modeled directed interactions (effective connectivity) with DCM.

Evaluating Models of Directed Influence

In previous work, Koechlin et al. (2003) used structural equation modeling on the time series of LPFC regions to provide evidence that episodic control increases effective connectivity from rostral-to-mid areas, while contextual control increases effective connectivity from mid-to-caudal areas, thereby producing a rostral-to-caudal cascade as a function of control demands. This Cascade Model was later updated to include the FPl at its apex (Koechlin and Summerfield, 2007). However, estimating effective connectivity directly on hemodynamic signals is problematic since what appears to be effective connectivity may arise due to hemodynamic confounds (David et al. 2008). Hence, accounting for the translation of neuronal signals to hemodynamic signals is preferred. Nee & D'Esposito (2016) used DCM to do just that and found that rostral-to-mid control influences and caudal-to-mid stimulus domain influences converged in mid-lateral areas. Amidst these convergent influences, the MFG showed the strongest outward directed influences suggesting that it may form the apex of hierarchical control. Nee & D'Esposito (2017) provided a modified model with many of the same properties. However, while both studies of Nee & D'Esposito included contextual control, neither included episodic control. Therefore, it remains unclear whether a mid-lateral convergence with MFG apex is a general control architecture, or whether control is better represented as a rostral-to-caudal cascade under other circumstances.

To adjudicate among these possibilities, we compared ten different models of LPFC function. Each model contained the same underlying fixed connectivity structure but varied in its conditional modulations and inputs. These models included 1) a null model with no modulations of effective connectivity by experimental manipulations, 2) a Cascade Model wherein episodic control modulated rostral-to-mid connectivity, and contextual control modulated mid-to-caudal connectivity, 3) a modified Cascade Model that added caudal-to-mid modulations as a function of stimulus domain. This model would provide a hybrid between the Cascade Model as original proposed, and the convergence observed in the models of Nee & D'Esposito. 4) The model of Nee & D'Esposito (2016) and 5) the model of Nee & D'Esposito (2017). In addition to varying modulations of effective connectivity by experimental manipulations, we entertained two sources of task inputs into the LPFC.

The first set of models posited that task inputs arrive at the MFG as assumed by Nee & D'Esposito (2016; 2017) and consistent with its role in representing task rules (Nee & Brown, 2012; Bunge et al., 2005; Genovesio et al., 2005; Wallis et al., 2001). A second set of models posited that task inputs arrive at the FPl potentially accounting for the increased control-related signals observed there as detailed in the ROI analyses above.

We employed a nested strategy to select amongst models. First, the five models with MFG task inputs were grouped as a family, and the five models with FPl task inputs were grouped as a family (Penny et al., 2010). Random-effects Bayesian model selection (BMS) was used to select among these families. Models that contained MFG cue inputs were superior to models that contained FPl cue inputs (exceedance probabilities: MFG family= 0.7261; FPl family= 0.2739) consistent with the assumptions of Nee & D'Esposito (2016).

Next, we grouped the remaining models by theoretical classes. Models were either grouped by their adherence to the null model (model 1), Cascade Model (models 2 and 3), or Nee & D'Esposito models (models 4 and 5). The Nee & D'Esposito family greatly outperformed the other families (exceedance probabilities: null= 0, cascade= 0.0004, Nee & D'Esposito= 0.9996). Directly comparing the five models revealed that model 5 (Nee & D'Esposito 2017's model) was overwhelmingly chosen in favor of the other four models (exceedance probabilities: model 5= 0.9971). We then utilized Parametric Empirical Bayes (PEB) at the subject and group levels to aggregate our parameter estimates while accounting for inter-subject variability. Inference on connections revealed that the MFG was the source of the strongest magnitude modulations of effective connectivity (Fig 7A). Moreover, examining fixed connectivity (Fig 7B) revealed a tendency for the MFG to send out more influences that it receives, consistent with an apical role.

To quantify this pattern, we contrasted the average magnitude of outward versus inward influences on fixed connectivity for each region. We refer to this quantity as hierarchical strength. Replicating previous work (Nee & D'Esposito, 2016; 2017), hierarchical strength peaked in the MFG ($t(24)= 3.4781$, $p_{corrected} = 0.0117$). The hierarchical strength of FPl was non-significantly negative (FPl: $t(24)= -1.6195$, $p_{corrected}=0.7104$) inconsistent with an apical role. Although non-significant, the sign of this effect was consistent with previous work (Nee & D'Esposito, 2016; 2017) (Fig 8). Collectively, these data replicate the findings of Nee & D'Esposito (2016; 2017).

Next, we directly compared the same model applied to the datasets of Nee & D'Esposito (2016; 2017). Changes in procedures lead to modest changes in parameters relative to those previously reported (Fig S2). When comparing across studies, fixed connectivity was remarkably consistent. The only significant difference observed across studies was that the fixed connection from the MFG upon the IFS was greater in the current dataset when compared to both the 2016 and 2017 samples (2016: $t(47)= 6.105$, $p_{corrected}<0.001$; 2017: $t(48)= 6.481$, $p_{corrected}<0.001$). Such data suggest that the general architecture of the LPFC is largely consistent across demands.

Examining changes in modulations, the only differences observed across studies involved modulations by the highest-level control process (episodic/temporal control) which were

reduced in the present dataset (relative to 2016: MFG to SFS, $t(47)= 3.311$, $p_{corrected}= 0.020$ and SFS to MFG, $t(48)= 2.673$, $p_{corrected}= 0.038$; relative to 2017: MFG to SFS, $t(47)= 4.271$, $p_{corrected}< 0.001$ and SFS to MFG, $t(48)= 3.007$, $p_{corrected}= 0.007$). These differences substantiate that temporal control and episodic control are distinct control processes, which modulate effective connectivity within the LPFC in dissociable ways. Such data suggest that the LPFC adapts to distinct control demands via task-specific modulations of a general architecture.

DISCUSSION

In the current study, we used a modified comprehensive control paradigm to test the generalizability of previous models of LPFC function. Relative to our past work (Nee & D'Esposito 2016; 2017), we altered the highest-level control process, and the dependency of lower-level control processes on higher-level control processes. These modifications allowed us to test whether our previously observed control architecture generalizes across diverse demands, or whether control architectures adapt according to task demands (Duncan, 2001; Duncan, 2010; Duncan et al., 2020). Behavioral and univariate fMRI data showed significant effects of our main factors indicating that our modified paradigm successfully elicited our processes of interest. However, activation patterns of our highest-level control process, episodic control, did not resemble our previous highest level control process, temporal control (Nee & D'Esposito 2016; 2017), and instead resembled contextual control. Nevertheless, contrasting control architectures using DCM provided evidence that our previously observed control architecture generalizes. Moreover, by comparing asymmetries of influence that denote hierarchical rank, we replicated the finding of an apical site in the mid-LPFC. This directed architecture was modulated by control demands with control demands that were common across studies (e.g. contextual control) leading to similar modulations across studies, while control demands that differed across studies (e.g. temporal vs episodic control) lead to distinct modulations. Such data indicate that the LPFC flexibly reconfigures to adapt to control demands. Collectively, these data indicate that the control architecture of the LPFC is both general and adaptive.

The present data speak to tensions regarding the fundamental organization of the LPFC. Although a number of studies have suggested that the LPFC is organized by abstraction of control demands (Koechlin et al., 2003; Nee & Brown, 2012; Bahlmann et al., 2014; Nee & D'Esposito, 2016; see Badre & Nee, 2018 for a review), others have failed to find a systematic organization (Reynolds et al., 2012; Crittenden and Duncan, 2014; Pischedda et al., 2017). Here, despite finding robust behavioral and neural signatures, we found no systematic differences between two cognitive control processes that putatively differ by abstraction: contextual control and episodic control. This contrasts starkly with our previous observations contrasting contextual control and temporal control that did lead to systematic differences. Interestingly, both temporal control in previous work and episodic control in the current paradigm operate over equivalent timeframes and can thus be considered to be equivalent as a function of temporal abstraction (Fuster, 2001; Badre, 2008; Badre and Nee, 2018). Therefore, the rostral-caudal axis of control across the PFC does not seem to be solely informed by temporal abstraction (see also Reynolds et al., 2012; Nee et al., 2014). Such data indicate the need to properly operationalize abstraction so as to determine the

situations that do and do not produce systematic differences across the rostral-caudal axis of the LPFC (see more below).

Despite marked differences in activations patterns observed at the putative highest-level of control when contrasting the present study with our previous work, we found that the same model architecture favored previously (Nee & D'Esposito, 2017) was also favored here. Moreover, fixed effective connectivity, reflecting the general directed architecture of the LPFC, was remarkably consistent across our studies with the lone replicable difference being an increase in MFG influence over IFS here. Critically, we replicated the finding that, insofar as hierarchy is reflected in dominance relationships, an area of the mid-LPFC is the apex of a putative LPFC control hierarchy. The generalizability of the directed architecture of the LPFC is reminiscent of observations that functional connectivity remains relatively fixed across diverse demands (Cole et al., 2014; Gratton et al., 2018). Such data suggest that the LPFC has a control architecture that generalizes across diverse tasks with an apex in the mid-LPFC.

On the backdrop of this general directed architecture, we observed that the LPFC adapts to distinct control demands via modulations of effective connectivity. In both the present and our past tasks, these modulations tended to coalesce towards the middle of the LPFC (i.e. in cMFG and IFS) potentially reflecting the integration of information to facilitate context-appropriate behavior (Nee, 2021; see also Cole et al., 2013). While we observed similar contextual control related modulations here as in past work, modulations as a result of our "highest" control demands differed across studies. Here, we observed no significant modulations of effective connectivity as a result of episodic control. Coupled with the observation that episodic control activations were largely overlapping with those elicited by contextual control suggests that in the present paradigm, there appears to be nothing that distinguishes episodic control over-and-above contextual control. By contrast, both activations and effective connectivity modulations distinguished temporal control from contextual control in previous work (Nee & D'Esposito, 2016; 2017). Hence, distinct control demands are associated with distinct modulations of effective connectivity.

Although we found a common model architecture best described both the present dataset and our previous data (Nee & D'Esposito, 2017), it is notable that the tasks used in both studies have a number of similarities. The most marked difference among the tasks regarded replacing temporal control with episodic control, which coincidentally also produced differences in effective connectivity estimates. Hence, one might question whether the similarities among model architectures reflects the similarities in the tasks. Conversely, one might wonder whether a markedly different task would produce markedly different effective connectivity estimates. A satisfying answer to this question would require examining the same control processes in a different task. To provide comparability to the present work, such a task would need to have sufficient power and specificity at the individual level to properly identify the areas of interest along the rostral-caudal and dorsal-ventral axes of the LPFC. Given the non-triviality of this need, we chose not to stray too far from the original task design. Nevertheless, novel task designs warrant exploration to more fully determine the extent to which the LPFC architecture we have described generalizes.

In the current dataset, episodic and contextual control shared largely overlapping activations. This is in contrast to distinctions between episodic and contextual control originally observed by Koechlin et al. (2003) and replicated by others (Bahlmann et al., 2014; Bahlman et al., 2015). It is notable that the contextual control activations observed here encompass both the episodic and contextual control areas observed by Koechlin and others. It is possible that by embedding contextual control demands as a subtask within a larger block, the subtask itself was regarded as an episode, thus engaging both episodic and contextual control. Nevertheless, a singular thread may run through both nominal contributions. Contextual control embodies the maintenance and selection of task sets while episodic control requires the maintenance of temporal information that informs task sets. Therefore, both main factors of control in the current paradigm share an emphasis on defining, selecting, and maintaining task sets. By contrast, temporal control as engaged by the paradigm of Nee & D'Esposito (2016; 2017) involved future-oriented planning that was not concerned with task sets. Future work that defines the representational nature of such plans as distinct from task sets may facilitate our ability to properly operationally define distinct forms of control.

Our definition of episodic control stems from the original terminology put forth by Koechlin et al. (2003). There, episodic control was defined as "the information conveyed by episodic signals (past events) and required for selecting task sets, when stimuli and contextual signals occur." The delay condition involved sustaining episodic signals, while the dual condition involved both sustaining episodic signals and using them to select task sets (i.e. to inform contextual control). In both cases, similar activation patterns were observed extending from caudal to rostral LPFC. Episodic control is also necessary in sequential tasks when the appropriate task to select depends upon position in a sequence (Badre and Nee, 2018). LPFC areas, including rostral LPFC, have been demonstrated to show ramping dynamics during such tasks (Desrochers et al., 2015; Wen et al., 2020). However, although tasks that require the selection of appropriate task sets as a function of sequential position involve episodic control, they also involve sequence monitoring processes to track ordinal position. Interestingly, sequence monitoring without the need to select task sets produces ramping dynamics in the same areas (Desrochers et al., 2019). Moreover, holding an item in mind for future integration produces similar ramping activations in rostral LPFC (De Pisapia et al., 2007; De Pisapia and Braver, 2008; Nee et al., 2014). Collectively, these data suggest that while episodic control is sufficient to activate the rostral LPFC, episodic control is not necessary to activate the rostral LPFC, and the rostral LPFC is not exclusive to episodic control.

The functions of the rostral-most area of the LPFC, the FPl, remains an intriguing mystery. Koechlin and Summerfield (2007) suggested that the FPl is primarily concerned with branching control: holding an episode in a pending state (see also Koechlin et al., 1999). Several studies have suggested a possible role of the FPl in the integration of internal representations into ongoing cognitive operations (Braver & Bongiolatti, 2002; De Pisapia et al., 2007; Nee et al. 2014). The FPl has also been implicated in the processes of monitoring counterfactual outcomes for future decisions (Boorman et al. 2009; Boorman et al. 2011). Here, we observed additive activations in the FPl for episodic and contextual control. These activations may reflect a general load of future-oriented representations, which may

facilitate switching amongst tasks and episodes. Interestingly, both fixed and modulatory influences of the FPl were negative. The negative influences might facilitate selection potentially by driving down competition and/or by sharpening the tuning of representations (Higo et al., 2011; Miller et al., 2011; Lee & D'Esposito, 2012). Alternatively, such negative influences may help to segregate future-oriented cognition from ongoing demands (Nee, 2021). Consistent with the latter, we have observed that disruption of the FPl by transcranial magnetic stimulation leads to errors during subtask phases involving temporal control (Nee & D'Esposito, 2017), potentially reflecting a premature intrusion of future-oriented cognition on current behavior. Collectively, these data suggest that the FPl is critical in organizing plans for the future.

Although the MFG did not show the strongest control-related activations among the areas examined, it is interesting to note that the under-additive activation pattern observed in this area matched the under-additive behavioral pattern observed on return trials. A similar pattern was also observed in the SFS, but not in the FPl. We speculate that the under-additive patterns in both activations and behavior reflect the improved integration of episodic control and contextual control through repeated use during the dual condition. However, why this leads to strong, additive control signals in the FPl and weaker, under-additive signals in the MFG and SFS remains unclear. Causal techniques that contrast the impact of direct perturbations of the MFG, FPl, and SFS may facilitate elucidating more specific, dissociable roles of these areas.

Here, as in past work, we have used functional activations to localize the area in mid-LPFC that appears apical. It may be possible to use anatomical landmarks in place of, or in addition to such procedures. It has been suggested that the area in mid-LPFC that we have referred to as "MFG" may correspond to a tertiary sulcus of the middle frontal gyrus: the posterior middle frontal sulcus (Miller et al., 2021). Individual variability in the location of such tertiary sulci may give rise to the observed variability in LPFC functional topography (Braga & Buckner, 2017; Gratton et al., 2018). Moreover, tertiary sulci are not well identified in group averaged atlases since their individual-level variability is not easily preserved in a group average (Miller et al., 2021). Future work examining the individual-level correspondence between the posterior middle frontal sulcus and a LPFC apex could facilitate more widespread examination of the role of this putatively apical area across other tasks, and whether it is altered in neurocognitive dysfunction.

To some extent, that the mid-LPFC (MFG) was observed to have the greatest influence among modeled areas may be surprising in the present work. Previously (Nee and D'Esposito, 2016; 2017), we observed that both the FPl and MFG increased activation in each of the conditions requiring additional cognitive control demands above baseline (i.e. dual, delay, restart) with both areas showing particularly more pronounced activations in the delay condition (requiring temporal control) relative to other areas. That the MFG appeared apical indicated that among the MFG and FPl, the MFG was the main driver of activation dynamics. However, in the present work, univariate activations in the MFG did not appear to stand out relative to other areas. Even still, the MFG was observed to have the largest influence among the areas modeled. This underscores the importance of effective connectivity modeling techniques such as DCM. It is important to note that

statistical modeling of univariate signals treats areas as independent and does not indicate how interactions among areas gives rise to the observed time series. Despite marked differences in the univariate signals in the MFG here relative to past work, DCM uncovered a similar pattern of influences in both fixed connectivity and modulations of connectivity by contextual control. Notably, fixed connectivity is based upon activations that do not vary as a function of condition, hence that the MFG is not modulated as a function of condition as much as some of the other areas is not necessarily discordant with it having strong fixed connectivity. This fixed connectivity may provide a scaffold for the modeled modulations. Such data suggest that strong univariate signaling need not be yoked to strong effective connectivity modulations. Interestingly, that the MFG shows marked influences here despite weaker univariate signaling suggests that small perturbations of the MFG can potentially give rise to large changes throughout the LPFC.

A general apex of cognitive control could be the target of intervention efforts and additional research. Deficits in cognitive control are a major symptom in many clinical populations (Soloman et al., 2009; Lesh et al., 2011; Goschke, 2014). Our data suggest that the MFG's widespread influence over control-related processing and behavior would mark it as a prime candidate for further directed investigation in control-deficit populations. However, there is evidence that suggests that the effects of noninvasive direct interventions in humans is varied across higher level cortex such as the PFC (Castrillon et al. 2020) and even across areas such as the primary motor cortex (M1) (Hamada et al. 2013). Further work investigating the specific effects of targeting the apical site suggested by this study is required in order to better understand how to best apply these findings.

Some limitations of the present work bear noting. DCM is a hypothesis-driven technique designed for model comparison. As such, model selection is only able to choose the best model from the options provided. While our model space was carefully constrained by previous research and neuroanatomical plausibility, it is by no means exhaustive. There is a distinct possibility that better fitting models of LPFC function are present within the vast number of potential combinations of conditional modulations. However, an exhaustive search of a model space this large is not computationally possible. Nevertheless, systematic exploration of other theoretically motivated models may reveal better fitting models for the current data set. These models should then be tested on past data to assess their own generalizability and elucidate more accurate models for future research.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGEMENTS

This work was supported by the National Institute of Mental Health (R01 MH121509 – Derek Nee) and features re-analyses of data supported by the National Institute of Mental Health (R01 MH063901 – Mark D'Esposito), and National Institute of Neurological Disorders and Stroke (P01 NS040813 – Mark D'Esposito; F32 NS0802069 – Derek Nee). The authors thank Chris Martin for insightful comments on this work.

REFERENCES

Andersson JL, Hutton C, Ashburner J, Turner R, & Friston K (2001). Modeling geometric deformations in EPI time series. Neuroimage, 13(5), 903–919. [PubMed: 11304086]

Ashburner J, & Friston K (1997). Multimodal image coregistration and partitioning—a unified framework. Neuroimage, 6(3), 209–217. [PubMed: 9344825]

Badre D, & D'Esposito M (2007). Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. Journal of cognitive neuroscience, 19(12), 2082–2099. [PubMed: 17892391]

Badre D (2008). Cognitive control, hierarchy, and the rostro–caudal organization of the frontal lobes. Trends in cognitive sciences, 12(5), 193–200. [PubMed: 18403252]

Badre D, D'Esposito M. Is the rostro-caudal axis of the frontal lobe hierarchical? Nat Rev Neurosci. 2009 Sep;10(9):659–69. doi: 10.1038/nrn2667. Epub 2009 Aug 12. . [PubMed: 19672274]

Badre D, & Nee DE (2018). Frontal cortex and the hierarchical control of behavior. Trends in cognitive sciences, 22(2), 170–188. [PubMed: 29229206]

Bahlmann J, Blumenfeld RS, & D'Esposito M (2015). The rostro-caudal axis of frontal cortex is sensitive to the domain of stimulus information. Cerebral cortex, 25(7), 1815–1826. [PubMed: 24451658]

Bahlmann J, Beckmann I, Kuhlemann I, Schweikard A, & Münte TF (2015). Transcranial magnetic stimulation reveals complex cognitive control representations in the rostral frontal cortex. Neuroscience, 300, 425–431. [PubMed: 26037799]

Bechara A, Damasio AR, Damasio H, & Anderson SW (1994). Insensitivity to future consequences following damage to human prefrontal cortex. Cognition, 50(1–3), 7–15. [PubMed: 8039375]

Boorman ED, Behrens TE, Woolrich MW, Rushworth MF. How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. Neuron. 2009 Jun 11;62(5):733–43. doi: 10.1016/j.neuron.2009.05.014. [PubMed: 19524531]

Boorman ED, Behrens TE, & Rushworth MF (2011). Counterfactual choice and learning in a neural network centered on human lateral frontopolar cortex. PLoS biology, 9(6), e1001093. [PubMed: 21738446]

Brainard DH (1997) The Psychophysics Toolbox, Spatial Vision 10:433–436. [PubMed: 9176952]

Braga RM, & Buckner RL (2017). Parallel interdigitated distributed networks within the individual estimated by intrinsic functional connectivity. Neuron, 95(2), 457–471. [PubMed: 28728026]

Braver TS, & Bongiolatti SR (2002). The role of frontopolar cortex in subgoal processing during working memory. Neuroimage, 15(3), 523–536. [PubMed: 11848695]

Bunge SA, Wendelken C, Badre D, & Wagner AD (2005). Analogical reasoning and prefrontal cortex: evidence for separable retrieval and integration mechanisms. Cerebral cortex, 15(3), 239–249. [PubMed: 15238433]

Castrillon G, Sollmann N, Kurcyus K, Razi A, Krieg SM, & Riedl V (2020). The physiological effects of noninvasive brain stimulation fundamentally differ across the human cortex. Science Advances, 6(5), eaay2739. [PubMed: 32064344]

Crittenden BM, & Duncan J (2014). Task difficulty manipulation reveals multiple demand activity but no frontal lobe hierarchy. Cerebral cortex, 24(2), 532–540. [PubMed: 23131804]

Cole MW, & Schneider W (2007). The cognitive control network: Integrated cortical regions with dissociable functions. Neuroimage, 37(1), 343–360. [PubMed: 17553704]

Cole MW, Reynolds JR, Power JD, Repovs G, Anticevic A, & Braver TS (2013). Multi-task connectivity reveals flexible hubs for adaptive task control. Nature neuroscience, 16(9), 1348–1355. [PubMed: 23892552]

Cole MW, Bassett DS, Power JD, Braver TS, & Petersen SE (2014). Intrinsic and task-evoked network architectures of the human brain. Neuron, 83(1), 238–251. [PubMed: 24991964]

David O, Guillemain I, Saillet S, Reyt S, Deransart C, Segebarth C, & Depaulis A (2008). Identifying neural drivers with functional MRI: an electrophysiological validation. PLoS biology, 6(12), e315. [PubMed: 19108604]

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. Cerebral cortex, 17(5), 993–1006. [PubMed: 16769743]

D'Mello AM, Gabrieli JDE, & Nee DE (2020). Evidence for hierarchical cognitive control in the human cerebellum. Current Biology, 30(10), 1881–1892. [PubMed: 32275880]

Dosenbach NU, Fair DA, Miezin FM, Cohen AL, Wenger KK, Dosenbach RA, … & Petersen SE. (2007). Distinct brain networks for adaptive and stable task control in humans. Proceedings of the National Academy of Sciences, 104(26), 11073–11078.

Duncan J (2001). An adaptive coding model of neural function in prefrontal cortex. Nature reviews neuroscience, 2(11), 820–829. [PubMed: 11715058]

Duncan J (2010). The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. Trends in cognitive sciences, 14(4), 172–179. [PubMed: 20171926]

Duncan J (2013). The structure of cognition: attentional episodes in mind and brain. Neuron, 80(1), 35–50. [PubMed: 24094101]

Duncan J, Assem M, & Shashidhara S (2020). Integrated intelligence from distributed brain activity. Trends in Cognitive Sciences, 24(10), 838–852. [PubMed: 32771330]

Friston KJ, Harrison L, & Penny W (2003). Dynamic causal modelling. Neuroimage, 19(4), 1273–1302. [PubMed: 12948688]

Genovesio A, Brasted PJ, Mitz AR, & Wise SP (2005). Prefrontal cortex activity related to abstract response strategies. Neuron, 47(2), 307–320. [PubMed: 16039571]

Goldman-Rakic PS. 1987. Circuitry of the primate prefrontal cortex and the regulation of behavior by representational memory. Plum F, Mountcastle VB (Eds). Handbook of Physiology. Bethseda: American Physiological Society, Vol 5. p 373–417.

Gratton C, Laumann TO, Nielsen AN, Greene DJ, Gordon EM, Gilmore AW, … & Petersen SE. (2018). Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. Neuron, 98(2), 439–452. [PubMed: 29673485]

Goschke T (2014). Dysfunctions of decision-making and cognitive control as transdiagnostic mechanisms of mental disorders: advances, gaps, and needs in current research. International journal of methods in psychiatric research, 23(S1), 41–57. [PubMed: 24375535]

Goulas A, Bastiani M, Bezgin G, Uylings HB, Roebroeck A, & Stiers P (2014). Comparative analysis of the macroscale structural connectivity in the macaque and human brain. PLoS computational biology, 10(3), e1003529. [PubMed: 24676052]

Hamada M, Murase N, Hasan A, Balaratnam M, & Rothwell JC (2013). The role of interneuron networks in driving human motor cortical plasticity. Cerebral cortex, 23(7), 1593–1605. [PubMed: 22661405]

Higo T, Mars RB, Boorman ED, Buch ER, & Rushworth MF (2011). Distributed and causal influence of frontal operculum in task control. Proceedings of the National Academy of Sciences, 108(10), 4230–4235.

Fuster JM. The prefrontal cortex--an update: time is of the essence. Neuron. 2001 May;30(2):319–33. doi: 10.1016/s0896-6273(01)00285-9. [PubMed: 11394996]

Kleiner M, Brainard D, Pelli D, 2007, "What's new in Psychtoolbox-3?" Perception 36 ECVP Abstract Supplement.

Koechlin E, Basso G, Pietrini P, Panzer S, & Grafman J (1999). The role of the anterior prefrontal cortex in human cognition. Nature, 399(6732), 148–151. [PubMed: 10335843]

Koechlin E, Ody C, Kouneiher F. The architecture of cognitive control in the human prefrontal cortex. Science. 2003 Nov 14;302(5648):1181–5. doi: 10.1126/science.1088545. [PubMed: 14615530]

Koechlin E, Summerfield C. An information theoretical approach to prefrontal executive function. Trends Cogn Sci. 2007 Jun;11(6):229–35. doi: 10.1016/j.tics.2007.04.005. Epub 2007 May 1. [PubMed: 17475536]

Lee TG, & D'Esposito M (2012). The dynamic nature of top-down signals originating from prefrontal cortex: a combined fMRI–TMS study. Journal of Neuroscience, 32(44), 15458–15466. [PubMed: 23115183]

Lesh TA, Niendam TA, Minzenberg MJ, & Carter CS (2011). Cognitive control deficits in schizophrenia: mechanisms and meaning. Neuropsychopharmacology, 36(1), 316–338. [PubMed: 20844478]

Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. Annu Rev Neurosci. 2001;24:167–202. doi: 10.1146/annurev.neuro.24.1.167. [PubMed: 11283309]

Miller BT, Vytlacil J, Fegen D, Pradhan S, & D'Esposito M (2011). The prefrontal cortex modulates category selectivity in human extrastriate cortex. Journal of Cognitive Neuroscience, 23(1), 1–10. [PubMed: 20586702]

Miller JA, Voorhies WI, Lurie DJ, D'Esposito M, & Weiner KS (2021). Overlooked tertiary sulci serve as a meso-scale link between microstructural and functional properties of human lateral prefrontal cortex. Journal of Neuroscience, 41(10), 2229–2244. [PubMed: 33478989]

Nee DE, & Brown JW (2012). Rostral–caudal gradients of abstraction revealed by multivariate pattern analysis of working memory. Neuroimage, 63(3), 1285–1294. [PubMed: 22992491]

Nee DE, Brown JW, Askren MK, Berman MG, Demiralp E, Krawitz A, & Jonides J (2013). A meta-analysis of executive components of working memory. Cerebral cortex, 23(2), 264–282. [PubMed: 22314046]

Nee DE, Jahn A, & Brown JW (2014). Prefrontal cortex organization: dissociating effects of temporal abstraction, relational abstraction, and integration with fMRI. Cerebral Cortex, 24(9), 2377–2387. [PubMed: 23563962]

Nee DE, & D'Esposito M (2016). The hierarchical organization of the lateral prefrontal cortex. Elife, 5, e12112. [PubMed: 26999822]

Nee DE, & D'esposito M (2017). Causal evidence for lateral prefrontal cortex dynamics supporting cognitive control. Elife, 6, e28040. [PubMed: 28901287]

Nee DE (2019). fMRI replicability depends upon sufficient individual-level data. Communications biology, 2(1), 1–4. [PubMed: 30740537]

Nee DE (2021). Integrative frontal-parietal dynamics supporting cognitive control. elife, 10, e57244. [PubMed: 33650966]

Pelli DG (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies, Spatial Vision 10:437–442. [PubMed: 9176953]

Penny WD, Stephan KE, Daunizeau J, Rosa MJ, Friston KJ, Schofield TM, & Leff AP (2010). Comparing families of dynamic causal models. PLoS computational biology, 6(3), e1000709. [PubMed: 20300649]

Petrides M (2005). Lateral prefrontal cortex: architectonic and functional organization. Philosophical Transactions of the Royal Society B: Biological Sciences, 360(1456), 781–795.

Pischedda D, Görgen K, Haynes JD, & Reverberi C (2017). Neural representations of hierarchical rule sets: the human control system represents rules irrespective of the hierarchical level to which they belong. Journal of Neuroscience, 37(50), 12281–12296. [PubMed: 29114072]

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(2), e30284. [PubMed: 22355309]

Romanski LM (2004). Domain specificity in the primate prefrontal cortex. Cognitive, Affective, & Behavioral Neuroscience, 4(4), 421–429.

Stephan KE, Penny WD, Daunizeau J, Moran RJ, & Friston KJ (2009). Bayesian model selection for group studies. Neuroimage, 46(4), 1004–1017. [PubMed: 19306932]

Stephan KE, Penny WD, Moran RJ, den Ouden HE, Daunizeau J, & Friston KJ (2010). Ten simple rules for dynamic causal modeling. Neuroimage, 49(4), 3099–3109 [PubMed: 19914382]

Smith DM, Perez DC, Porter A, Dworetsky A, & Gratton C (2021). Light through the fog: using precision fMRI data to disentangle the neural substrates of cognitive control. Current Opinion in Behavioral Sciences, 40, 19–26. [PubMed: 33553511]

Solomon M, Ozonoff SJ, Ursu S, Ravizza S, Cummings N, Ly S, & Carter CS (2009). The neural substrates of cognitive control deficits in autism spectrum disorders. Neuropsychologia, 47(12), 2515–2526. [PubMed: 19410583]

Tanji J, & Hoshi E (2001). Behavioral planning in the prefrontal cortex. Current opinion in neurobiology, 11(2), 164–170. [PubMed: 11301235]

Tanji J, & Hoshi E (2008). Role of the lateral prefrontal cortex in executive behavioral control. Physiological reviews, 88(1), 37–57. [PubMed: 18195082]

Vincent JL, Kahn I, Snyder AZ, Raichle ME, & Buckner RL (2008). Evidence for a frontoparietal control system revealed by intrinsic functional connectivity. Journal of neurophysiology, 100(6), 3328–3342. [PubMed: 18799601]

Wallis JD, Anderson KC, & Miller EK (2001). Single neurons in prefrontal cortex encode abstract rules. Nature, 411(6840), 953–956. [PubMed: 11418860]

Zeidman P, Jafarian A, Corbin N, Seghier ML, Razi A, Price CJ, & Friston KJ (2019). A guide to group effective connectivity analysis, part 1: First level analysis with DCM for fMRI. Neuroimage, 200, 174–190. [PubMed: 31226497]

Zeidman P, Jafarian A, Seghier ML, Litvak V, Cagnan H, Price CJ, & Friston KJ (2019). A guide to group effective connectivity analysis, part 2: Second level analysis with PEB. Neuroimage, 200, 12–25. [PubMed: 31226492]

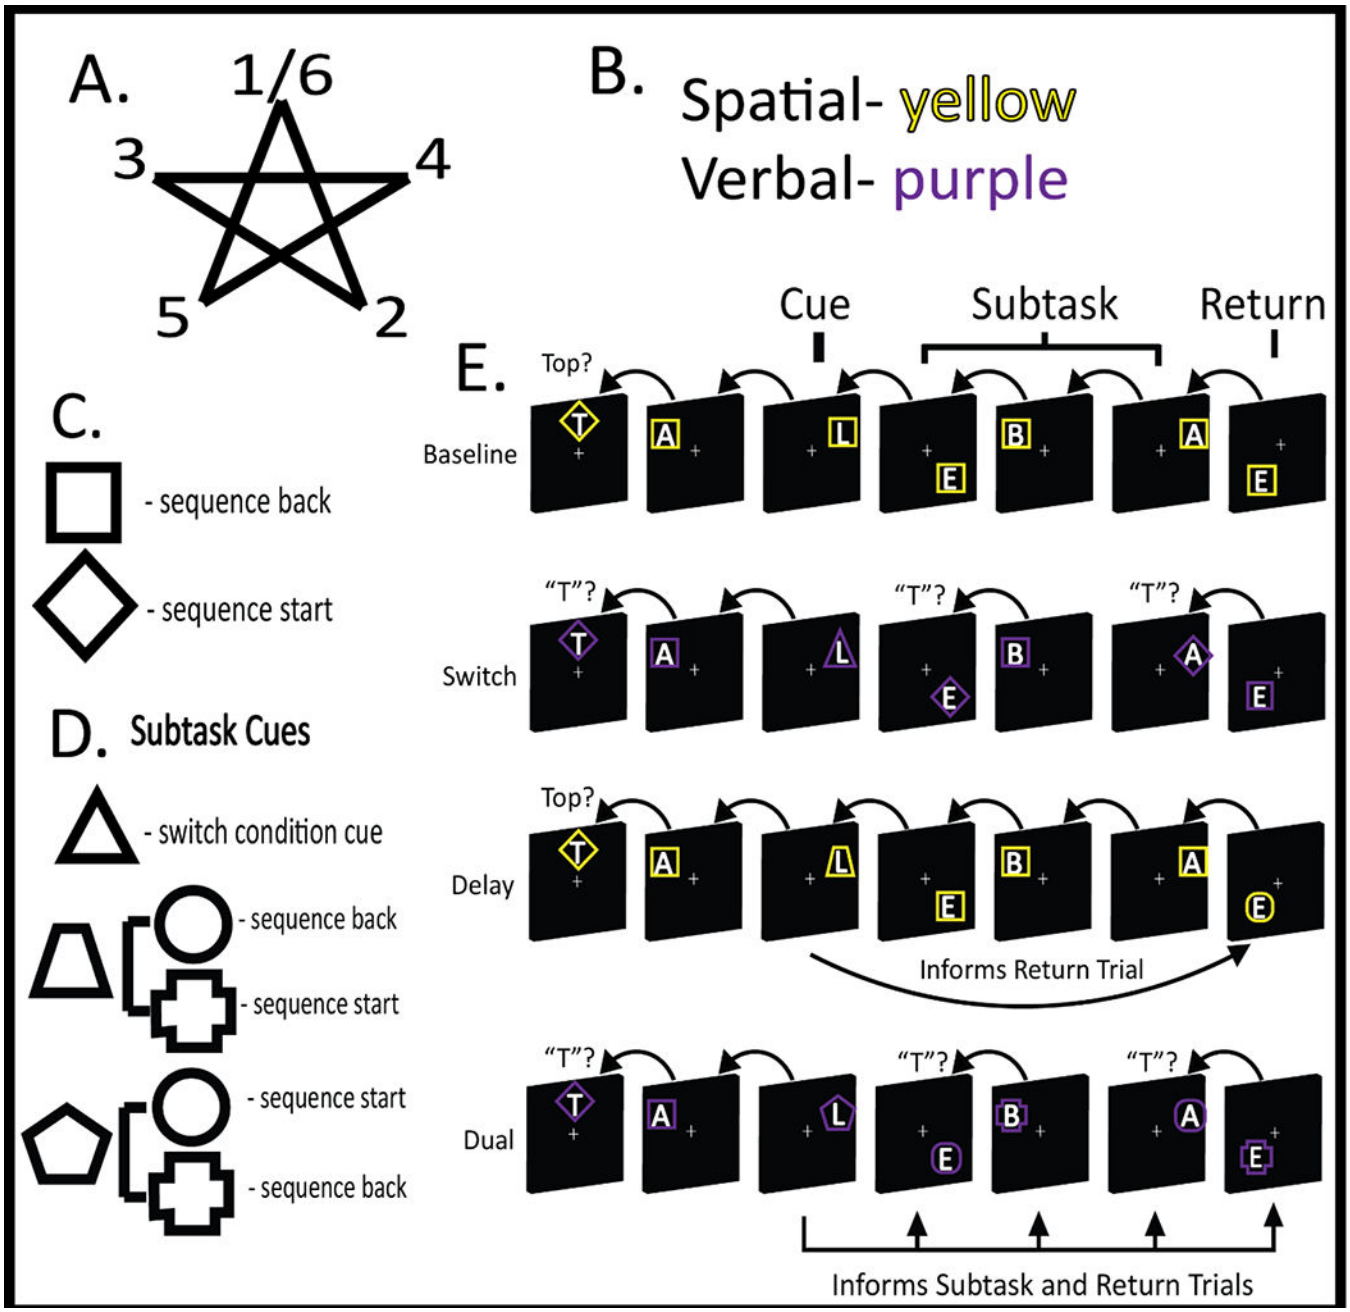

Fig 1:

Experimental paradigm. **A**) The spatial sequence participants used during spatial blocks of trials. The sequence began and ended at the top position (1/6). An analogous verbal sequence consisted of the ordering of the letters in the word "tablet." **B**) A potential color mapping indicating the attended domain within a block. The color of the frame surrounding the letter indicated whether locations (spatial; yellow) or letters (verbal; purple) were relevant for a given block. **C**) Frame shapes indicated the task: sequence back – determine whether the present stimulus follows the previous stimulus in the relevant sequence (square); sequence start – determine whether the present stimulus is the start of the relevant sequence.

D) Subtask cues indicated forthcoming control demands during a subtask phase. A switch condition cue (triangle) indicated that participants would switch among sequence start and sequence back tasks during the subtask. An episodic cue (trapezoid, pentagon) informed future shape-task mappings to be relevant at the end of the subtask phase (delay condition) or throughout the subtask phase (dual condition). **E**) Experimental trials by condition. Time unfolds left to right, backwards arrows indicate referents for sequence-back decisions, and forward arrows indicate the influence of cues upon task selection (episodic control). In the baseline condition, after an initial sequence start trial, all remaining trials required sequence-back. This required a constant engagement of sensorimotor control. The switch condition required switching between task sets during the subtask phase engaging contextual control. The delay condition required sustaining the shape-task mapping throughout the subtask phase to accurately respond to the return trial, thereby requiring episodic control. The dual condition required sustaining the shape-task mapping throughout the subtask phase (episodic control) and using it to select task sets (contextual control). Hence, in this condition, episodic control informs contextual control.

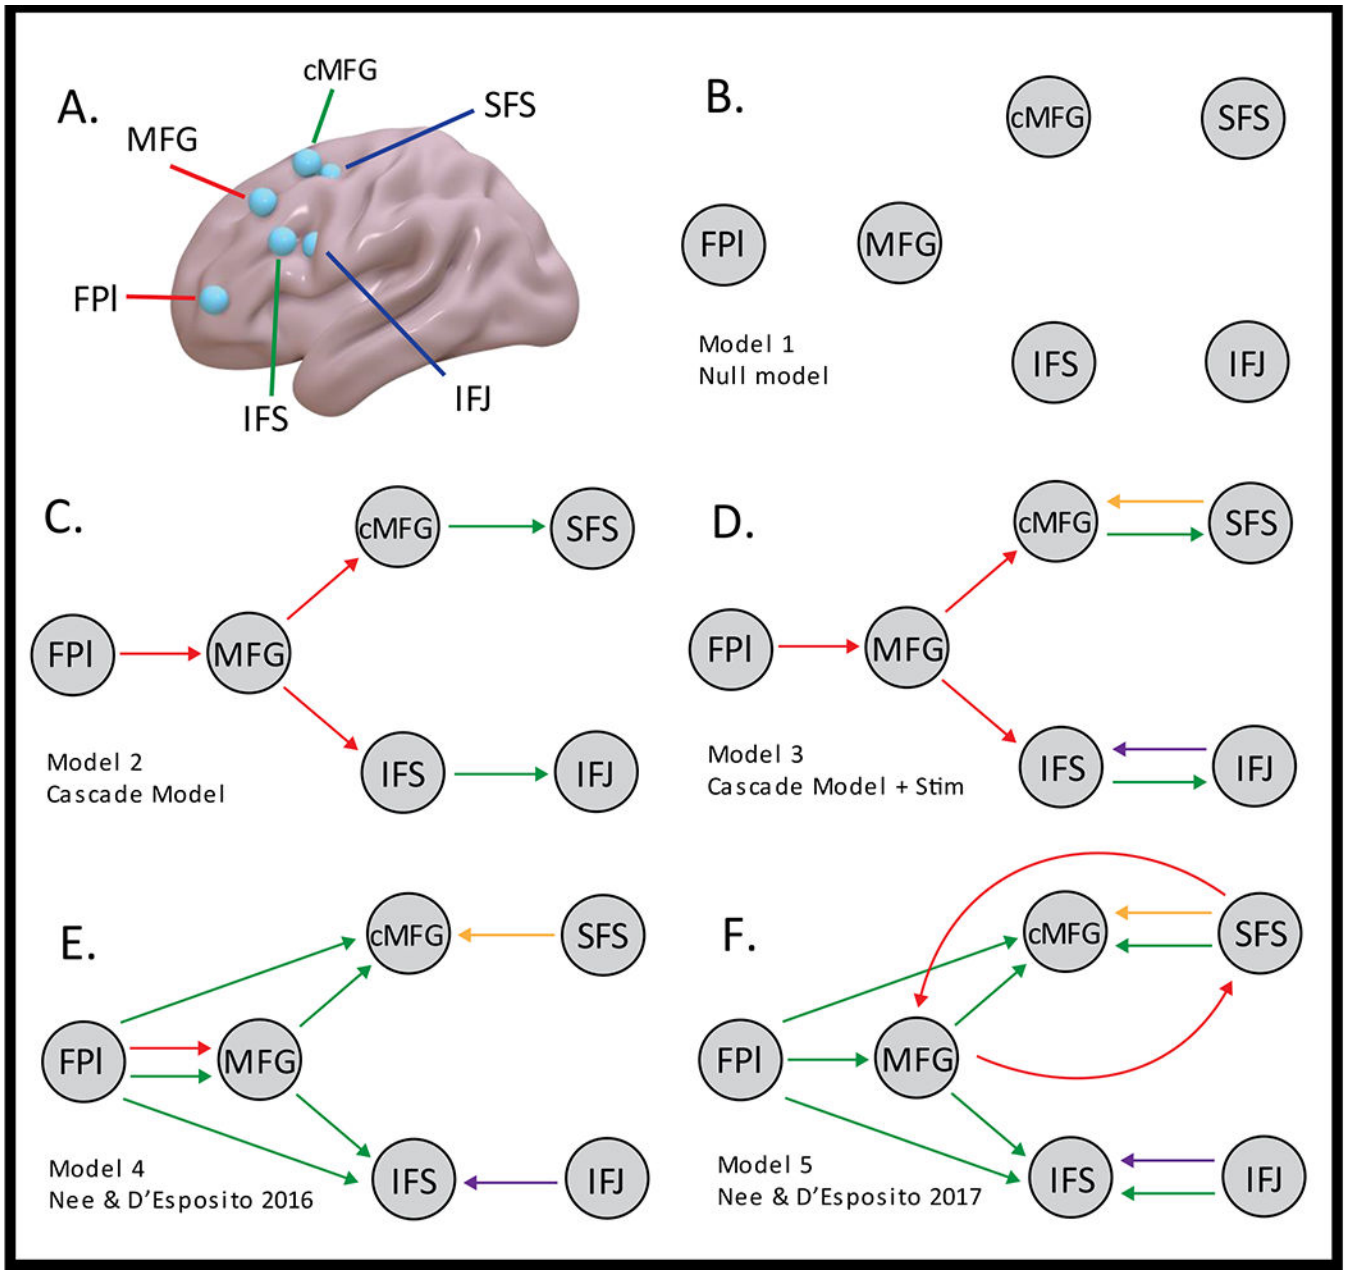

Fig 2:

ROIs and candidate models for DCM **A**) ROIs based upon Nee & D'Esposito 2016. **B-F**: Candidate models. The color of each arrow corresponds to the type of modulation (green = contextual control, red = episodic control, orange = spatial, purple = verbal). **B**) Null model without conditional modulations **C**) Theoretical Cascade model based upon predictions made in Koechlin et al., (2003) and Koechlin and Summerfield (2007). **D**) Cascade model with caudal to middle sensorimotor modulations. **E**) Nee & D'Esposito's 2016 model. **F**) Nee & D'Esposito's 2017 model. **F**) All 5 models were tested with cue inputs to MFG or FPl creating 10 total models.

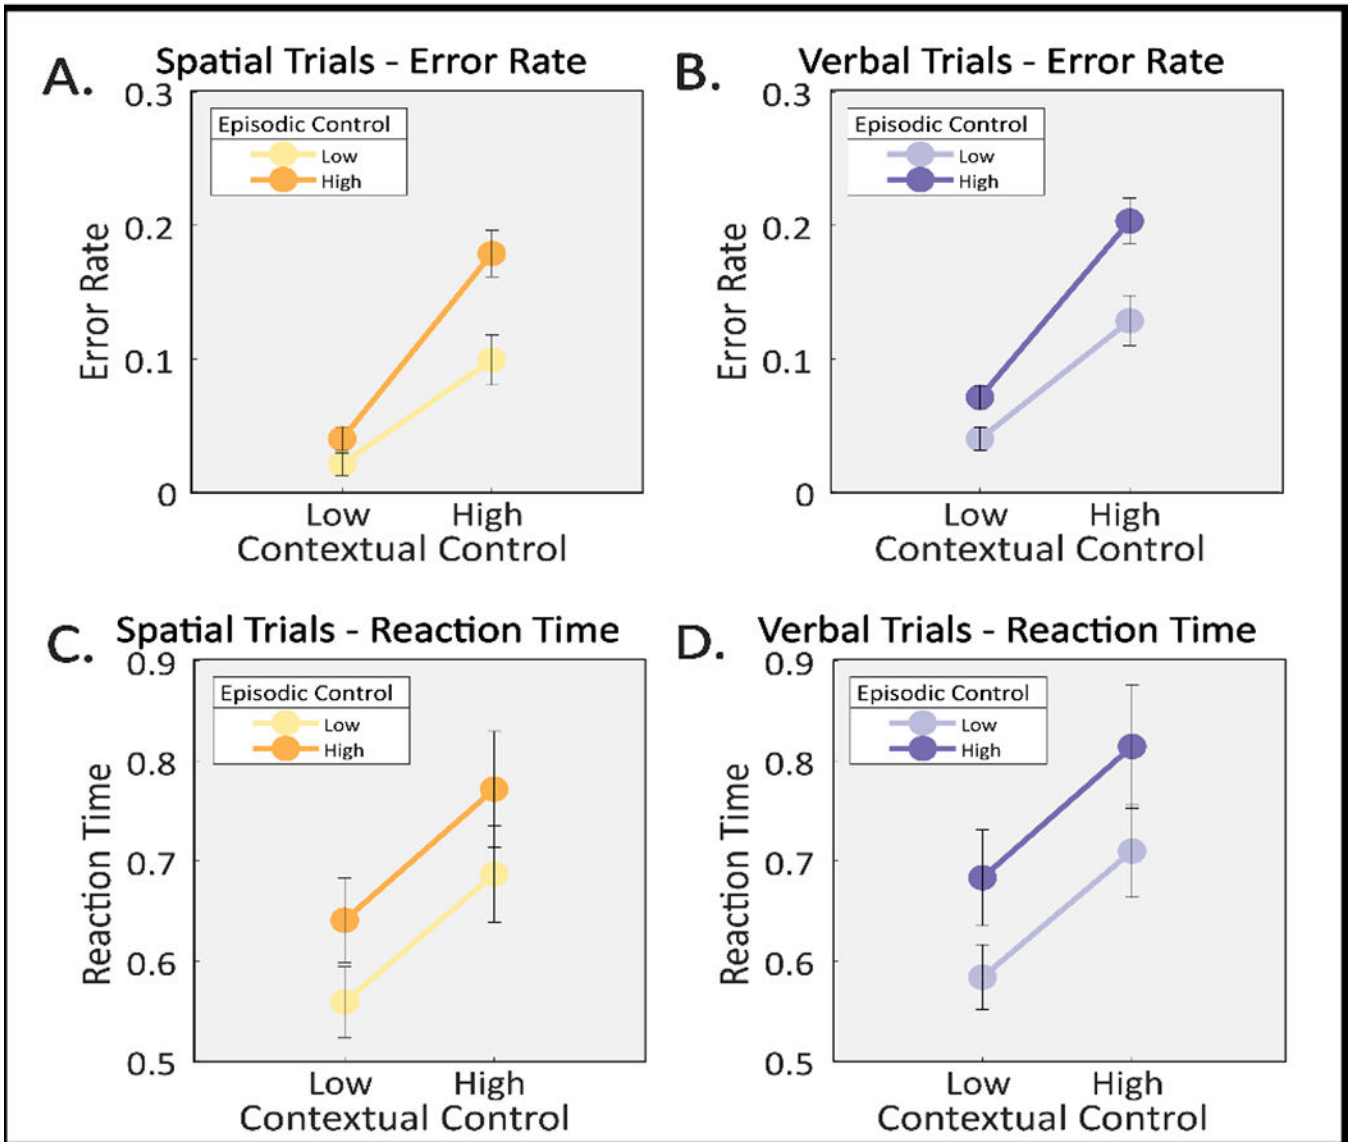

Fig 3:

Behavioral data on subtask trials. Error rates for spatial (A) and verbal (B) subtask trials plotted on the top showing main effects of episodic control, contextual control, and their interaction, as well as a main effect of stimulus domain. Reaction times for spatial (C) and verbal (D) subtask trials plotted on the bottom showing main effects of episodic control and contextual control. Spatial trials plotted in orange/yellow, verbal trials plotted in purple/lavender; darker colors denote high episodic control (orange; purple), lighter colors denote low episodic control (yellow; lavender).

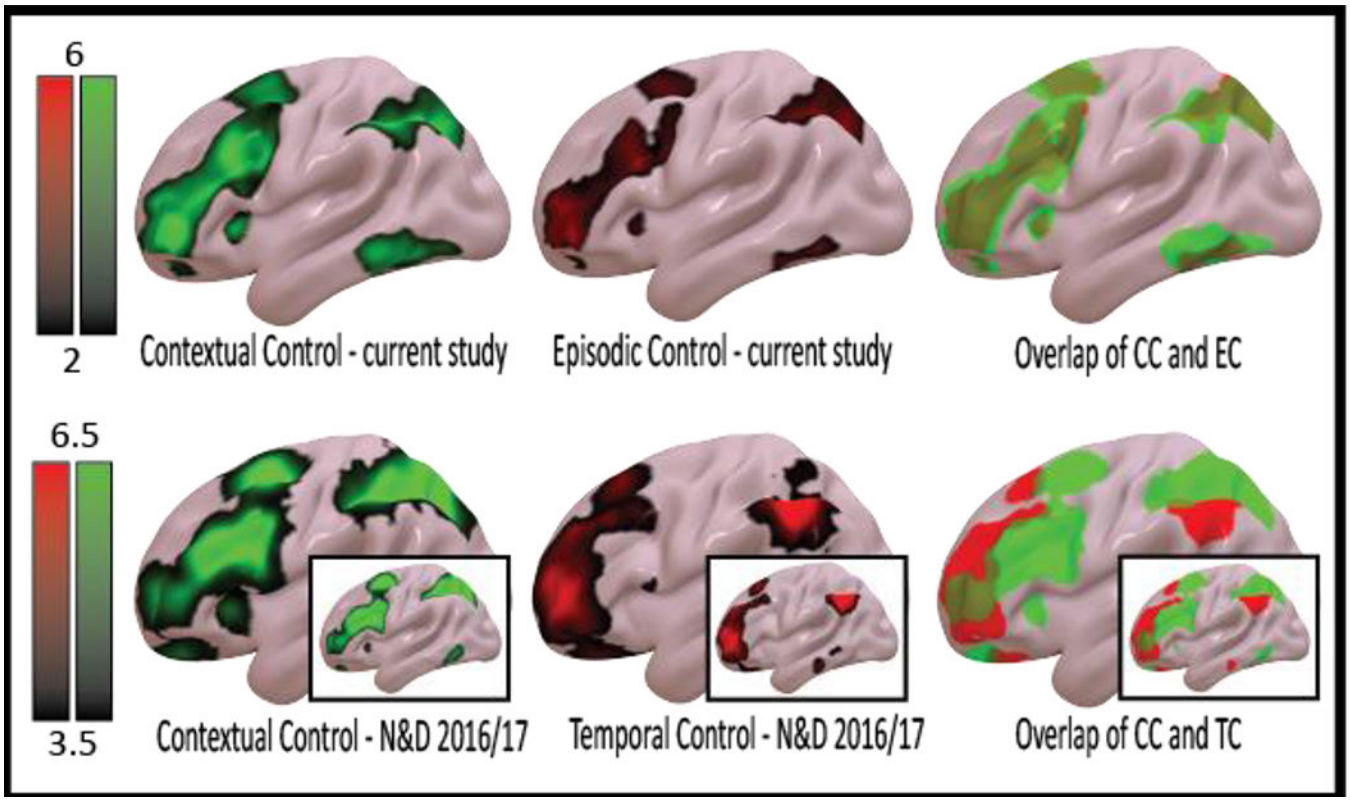

Fig 4:

Whole-brain control-related activations. **Top**: Activations for contextual control (dual + switch > delay + baseline; green), episodic control (dual + delay > switch + baseline; red), and their overlap. No clear topographical differentiation was observed between the contrasts. **Bottom**: Corresponding activations in Nee & D'Esposito (2016) for contextual control (green), temporal control (red), and their overlap. Insets display replicated contrasts from Nee & D'Esposito (2017). In contrast to the present study, the more abstract control contrast elicited more rostral activations in the LPFC. Contrasts for the present study depicted at t-statistics between 2-6. The datasets of Nee & D'Esposito are shown at t-statistics between 3.5-6.5 for visualization purposes.

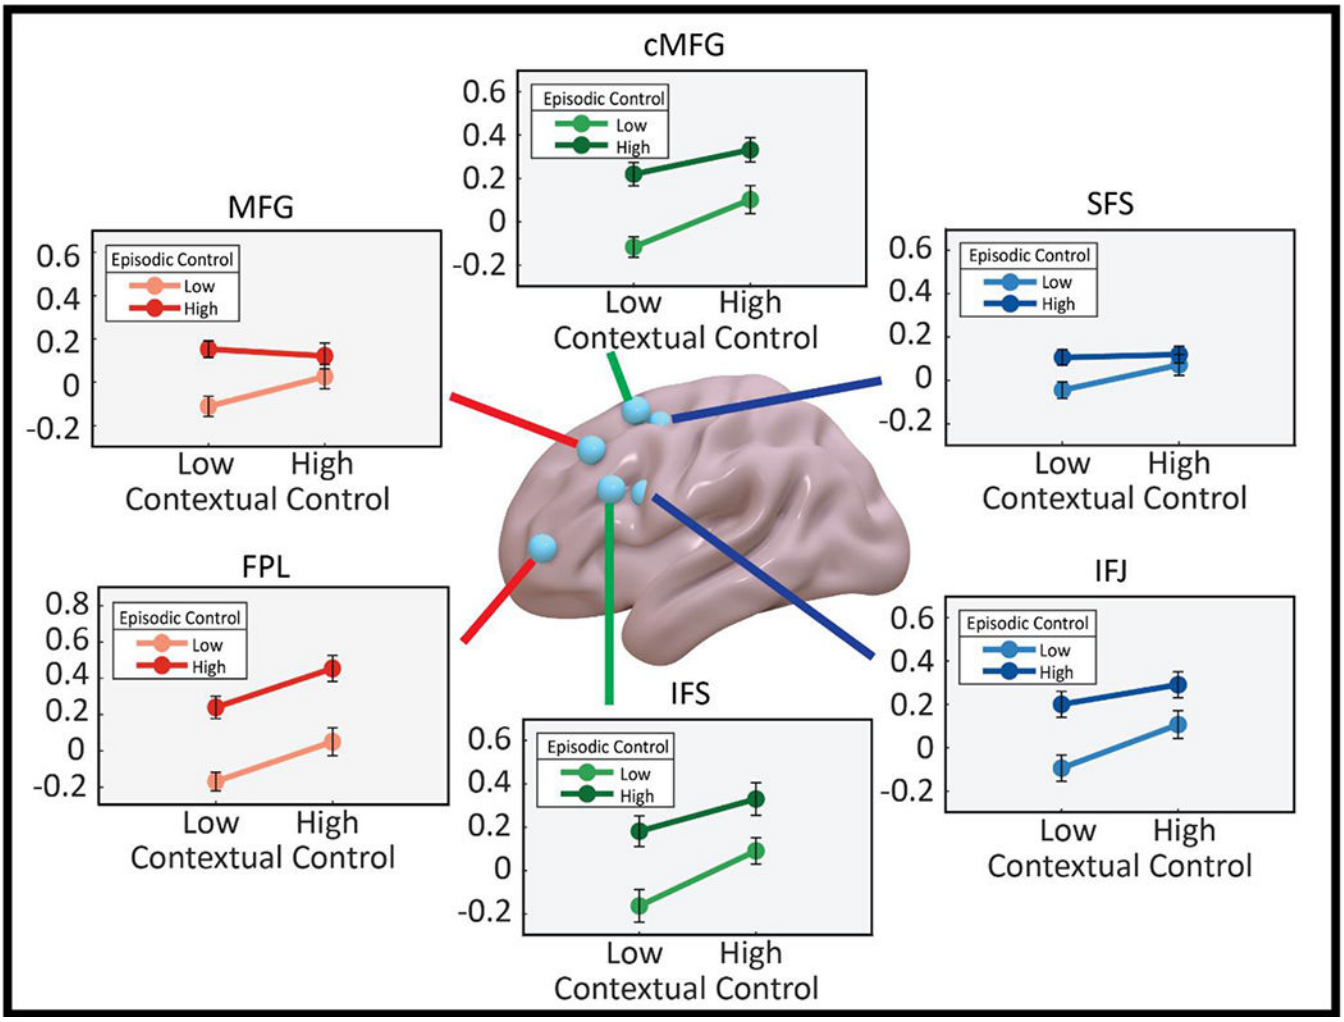

Fig 5:

ROI analyses. Condition related activations in each ROI for high and low contextual and episodic control. Overall, effects of contextual control were present in all ROIs. Significant effects of episodic control were observed in all areas except SFS and MFG. Both the SFS and MFG displayed an under additive interaction.

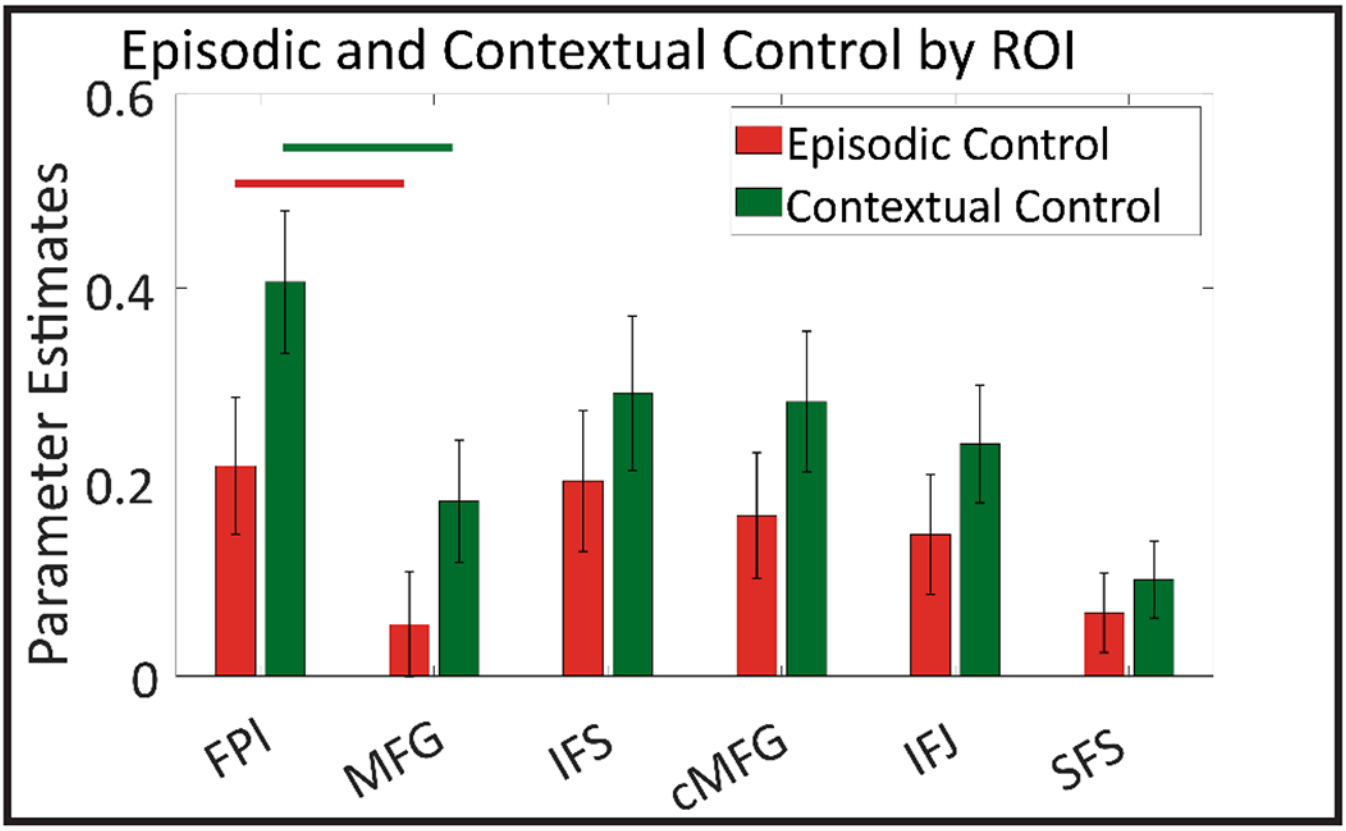

Fig 6:

Activation pattern across each ROI. Episodic control (dual+delay > switch+baseline; red) and Contextual control (dual+switch > delay+baseline; green). Control related activations elicited by our control demands tended to increase from caudal to rostral areas. However, MFG deviated from this pattern by showing reduced activations compared to FPl. Solid lines between FPl and MFG show significant differences at $p<0.001$ for episodic (red) and contextual control (green).

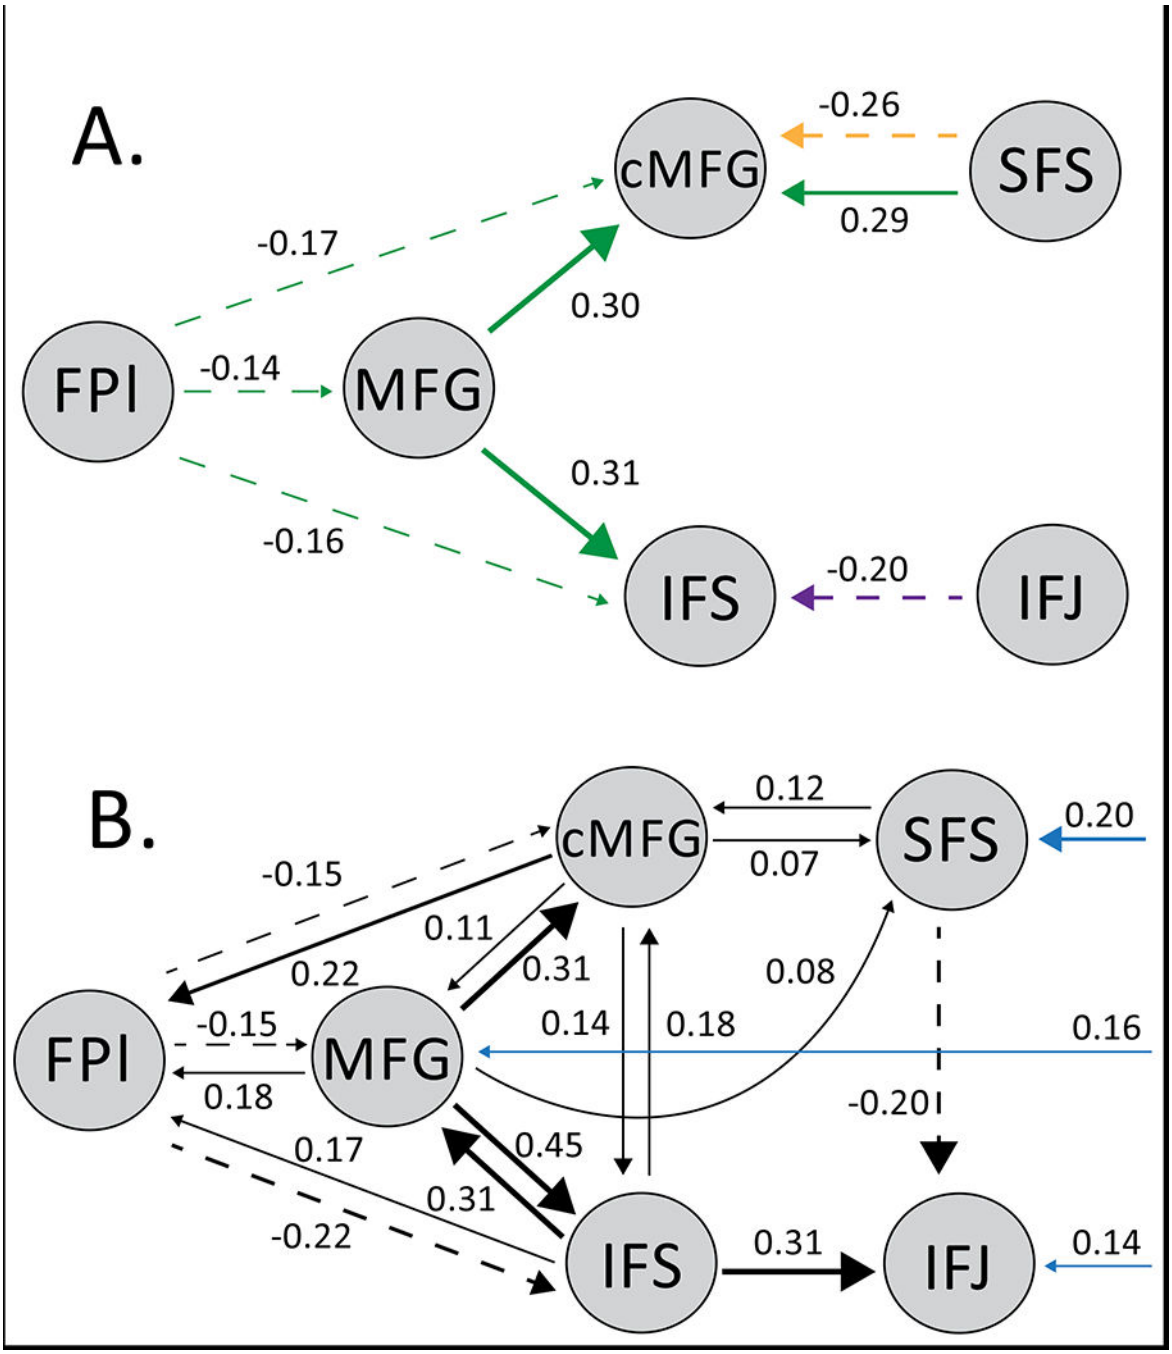

Fig 7:

Group level PEB showing significant parameter estimates. **A**) Modulations in connectivity by condition (green= contextual control, orange= spatial, purple= verbal). No modulations during episodic control passed significance. **B**) Fixed connectivity (black) and inputs (blue) to the model.

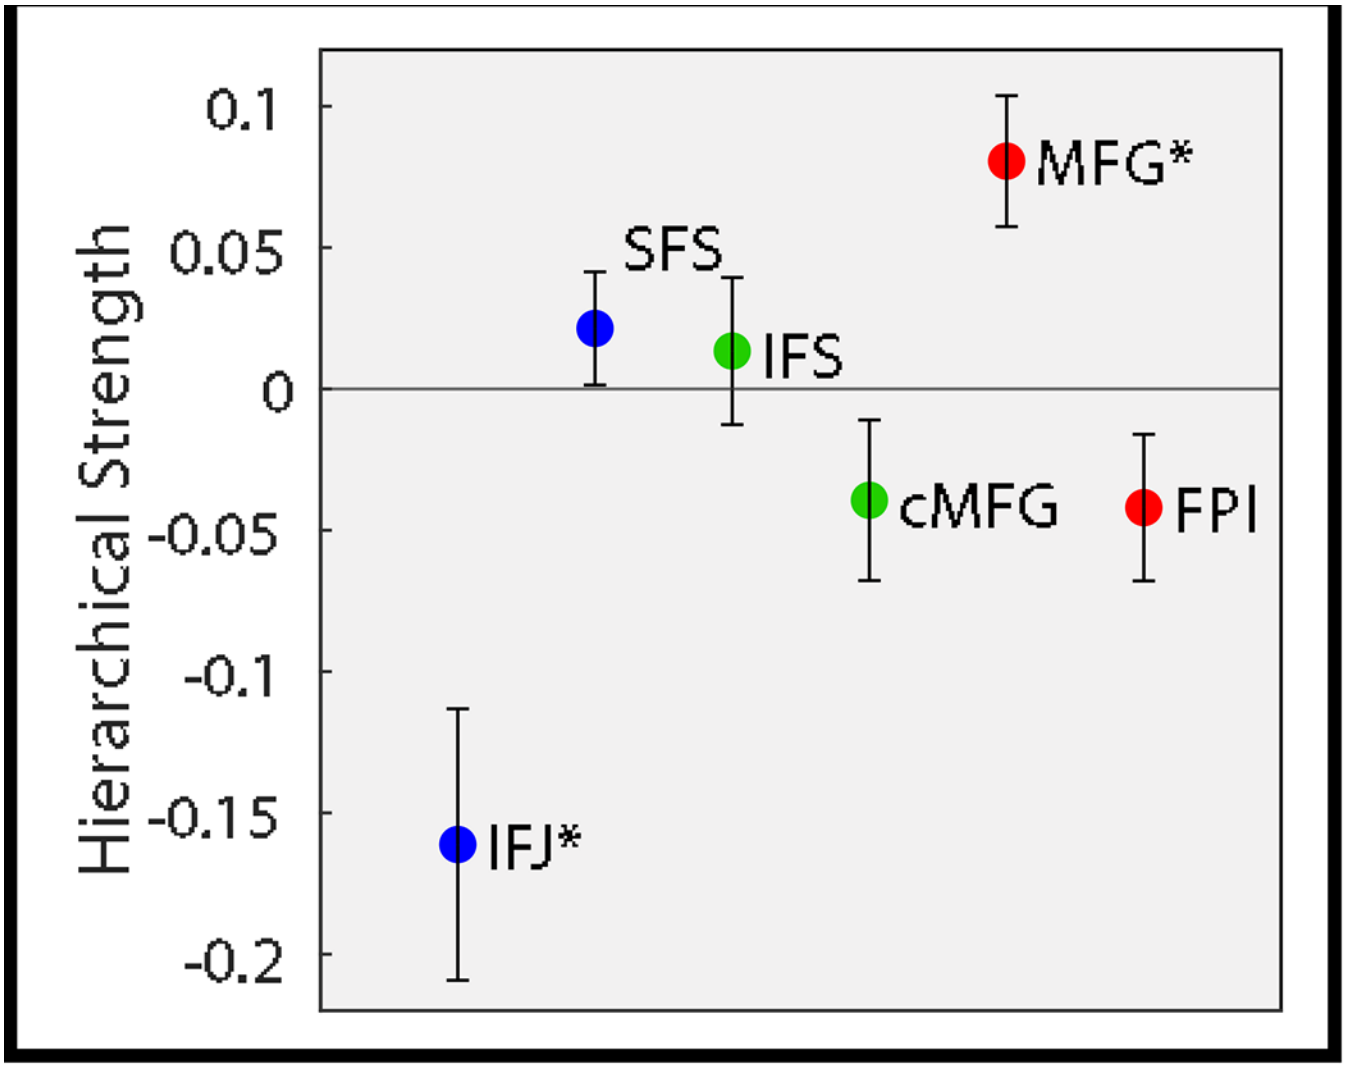

Fig 8:

Summarization of hierarchical strength of each region. Hierarchical strength was defined by the magnitude of outgoing connections minus incoming connections. Hierarchical strength peaked in the MFG and was non-significantly negative in FPl. * $p_{corrected}<0.05$.