

Published in final edited form as:

Nat Neurosci. 2009 April ; 12(4): 515–522. doi:10.1038/nn.2277.

Hierarchical cognitive control deficits following damage to the human frontal lobe

David Badre^{1,2}, Joshua Hoffman³, Jeffrey W Cooney³, and Mark D'Esposito^{3,4}

¹ Department of Cognitive and Linguistic Sciences, Brown University, Providence, Rhode Island, USA

² Department of Psychology, Brown University, Providence, Rhode Island, USA

³ Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley, California, USA

⁴ Department of Psychology, University of California, Berkeley, Berkeley, California, USA

Abstract

Cognitive control permits us to make decisions about abstract actions, such as whether to e-mail versus call a friend, and to select the concrete motor programs required to produce those actions, based on our goals and knowledge. The frontal lobes are necessary for cognitive control at all levels of abstraction. Recent neuroimaging data have motivated the hypothesis that the frontal lobes are organized hierarchically, such that control is supported in progressively caudal regions as decisions are made at more concrete levels of action. We found that frontal damage impaired action decisions at a level of abstraction that was dependent on lesion location (rostral lesions affected more abstract tasks, whereas caudal lesions affected more concrete tasks), in addition to impairing tasks requiring more, but not less, abstract action control. Moreover, two adjacent regions were distinguished on the basis of the level of control, consistent with previous functional magnetic resonance imaging results. These results provide direct evidence for a rostro-caudal hierarchical organization of the frontal lobes.

The function of the prefrontal cortex (PFC) is closely associated with cognitive control or the ability of humans and other primates to internally guide behavior in accordance with goals, plans and broader contextual knowledge^{1–18}. Consider the simple example of entering a colleague's office and finding a place to sit down. On a daily basis, in one's own office, the chair behind the desk is the appropriate seat. In another's office, however, we easily select the chair in front of the desk as being the socially appropriate choice. Overcoming a habitual tendency in order to coordinate behavior with an abstract social rule is an example of cognitive control.

From one perspective, cognitive control mechanisms operate through a process of biased competition, whereby maintenance of a distributed neural representation of the task context (colleague's office) configures processing throughout the action system to bias selection of an appropriate behavior (sit in the chair in front of the desk) over a competing one (sit in the chair behind the desk)^{8,19–21}. The frontal lobes are thought to be centrally involved in coding such contextual representations to provide internal control over action^{14,22,23}. However, the functional organization of the frontal lobe remains unknown. Although it is

Correspondence should be addressed to D.B. (david_badre@brown.edu).

Note: Supplementary information is available on the Nature Neuroscience website.

Reprints and permissions information is available online at <http://npg.nature.com/reprintsandpermissions/>

widely believed that separate frontal regions support distinct forms of control, there is little evidence to date of double dissociations in lateral frontal cortex and no evidence in human patients⁴. Thus, a fundamental goal in cognitive neuroscience is to characterize the functional organization of frontal cortex that supports the control of action.

Control of action can involve abstract goals, such as deciding whether to e-mail or call a friend, as well as the concrete motor programs required to carry out these abstract goals, such as selecting the appropriate sequence of keystrokes to type an e-mail greeting^{17,24–26}. Thus, computational models of cognitive control and of complex action have often included hierarchical architectures that represent such actions at different levels of abstraction^{24–27}. Consistent with the concept of hierarchically arrayed levels of control, neuroimaging studies have repeatedly demonstrated differences in functional activation along the rostral-caudal axis of lateral frontal cortex, ranging from dorsal premotor cortex (PMd; ~Brodmann area 6/8) to lateral frontal polar cortex (Brodmann area 10), such that more anterior regions were associated with progressively more abstract action control further removed from the selection of a concrete motor response^{28–31}. Across these previous studies, abstraction has been defined in different, although not necessarily mutually exclusive, ways³². Some have suggested that posterior regions are more sensitive to domain distinctions, such as spatial versus object, whereas more anterior regions are not^{18,33}. Others propose that progressively anterior regions coordinate action over longer time scales and so can maintain action representations and mediate action contingencies over longer temporal gaps^{7,34}. Still others have proposed that progressively anterior regions maintain more complex rules that choose a class of more specific, lower-level rules; the lowest being the rule that specifies a motor responses²⁸. For example, the choice to write an e-mail is abstract relative to choices about what words to put in the e-mail itself. Regardless of the specific definition of abstraction, the data consistently demonstrate that more rostral regions of frontal cortex are associated with progressively abstract control demands and representations.

Some theorists further interpret these data as reflecting a hierarchical organization of lateral frontal cortex, whereby control processes or representations at a given locus in the frontal lobes are influenced by more abstract control processing in ‘higher’, more anterior regions, but not in ‘lower’, more posterior regions. Such a hierarchical influence could reflect the passing/summing of control signals from anterior to posterior in the frontal lobe³¹, or the reduction of uncertainty at lower levels by action pathways chosen at higher levels²⁸ or by activating/coordinating task sets among lower-order processors^{35–37}. Anatomical evidence suggests that there is an asymmetry in the corticocortical connections in frontal cortex that could support such a processing hierarchy^{38,39}. Indirect evidence from effective connectivity analysis of functional magnetic resonance imaging (fMRI) data also supports an asymmetric anterior to posterior flow of influence^{31,35–37}. However, the neuroimaging data cannot be conclusive on this point. Indeed, some perspectives can account for a rostral-caudal functional gradient without a requirement that the processing architecture be hierarchical^{11,29,40,41}. Thus, a fundamental issue to resolve is whether the observed rostral-caudal gradient reflects a hierarchical or nonhierarchical organization of function³².

An anterior-to-posterior flow of control processing in the frontal lobes predicts that performance on tasks involving higher-order control should be impaired by disruptions to lower-order processors, even when the higher-order processors are intact. However, the reverse prediction should not hold. Performance should be unaffected for tasks involving only intact lower-order processors when higher-order processors are impaired. This hypothesized asymmetric pattern of deficit cannot be directly tested with neurophysiological methods, such as fMRI, electroencephalography or single-unit recording. Rather, it requires a lesion method that leads to isolated disruption of specific processors along the proposed hierarchical gradient.

To test the asymmetry hypothesis, we asked 12 individuals with focal frontal lobe lesions and 24 age-matched controls to perform a set of four response-selection tasks (Fig. 1a and Methods) that required increasing levels of hierarchically ordered control to select a correct key-press response. In other words, from the response to feature to dimension to context task (Fig. 1b), the appropriate representations to be selected increased in abstraction. To manipulate control in each of the four response-selection tasks, there were low-, mid- and high-competition conditions. These required either no selection (low) or selection from two (mid) or four (high) candidate representations at a given level of abstraction, respectively. For example, in the response task (first order of abstraction), participants selected a response on the basis of a learned mapping with a color cue presented on each trial. Competition increased as participants went from having no choice (one response, low) to having two (mid) or four (high) responses to choose from. As the hierarchical rank of the tasks increased from response to feature to dimension to context (Fig. 1b), the mid- and high-competition conditions of each task required selection of a more abstract representation than the task ranked below it, and so demanded higher-order control. For example, rather than only requiring selection of a response on the basis of a learned mapping (the response task), the mid- and high-competition conditions of the feature task (second-order control) required selection of a set of response mappings over other competitor sets. This logic was carried up four levels of abstraction across the four tasks.

In the low-selection condition of each task, competition was set equivalent to that of the mid-competition condition of the task ranked immediately below it (this logic is spelled out explicitly in Fig. 1b). For example, the low-competition condition of the feature task (Fig. 1b) required selection from among two responses, but from only one response set (defined on the basis of the target). Thus, the control demands for this condition were identical to the mid-competition condition of the response task, which also required selection from two responses, but from only one response set (Fig. 1b). In contrast, the mid- and high-competition conditions of the feature task required the selection of a response set from two or four candidate response sets, respectively. Thus, the low-competition condition provides an estimate of the contribution of lower-order control demands. Moreover, comparison of the mid- and high- with the low-competition condition provides within-task control for superficial differences between the tasks themselves.

Using fMRI, we previously demonstrated that the hierarchical level of control in these tasks determines the locus of activation along the rostro-caudal axis of lateral frontal cortex, with response, feature, dimension and context control being associated with PMd (~Brodmann area 6, Montreal Neurological Institute (MNI) standard: $x = -30, y = -10, z = 68$), anterior PMd/posterior PFC (pre-PMd, ~Brodmann area 8, MNI: $x = -38, y = 10, z = 34$), inferior frontal sulcus (IFS, ~Brodmann area 45, MNI: $x = -50, y = 26, z = 24$; Brodmann area 9/46, MNI: $x = -52, y = 28, z = 38$) and frontal polar cortex (~Brodmann area 10, MNI: $x = -36, y = 50, z = 6$), respectively²⁸. In contrast, the parametric increase in competition (low, mid and high) was associated with a corresponding increase in activation only at the frontal locus supporting that level of control²⁸.

Depending on the site of damage to the frontal lobe, individuals with lesions should be impaired for the mid- and high-competition conditions at the level of abstraction at which disruption of control has occurred and all conditions of tasks at more abstract levels, despite having intact control processors at these levels. In contrast, individuals with lesions should perform normally on the low-competition condition of the impaired level and all conditions of tasks at lower levels. In this study, we tested this hypothesis most directly for two hierarchical levels. Specifically, we tested that a lesion to pre-PMd should impair performance on the mid- and high-competition conditions of the feature task and all conditions of the dimension and context tasks because these all require a second order of

control. However, this lesion should not impair performance on any condition of the response task or the low condition of the feature task because these only require a first order of control. In contrast, a lesion to IFS should impair performance on the mid- and high-competition conditions of the dimension task and all conditions of the context task because these all require a third order of control. However, such a lesion should not impair performance on any condition of the response or feature task or the low condition of dimension task because these only require first and second orders of control. Such a pattern of results would be direct evidence for hierarchy in the frontal lobe.

RESULTS

In general, individuals with frontal lesions demonstrated increasingly poor performance as demands on control increased in abstraction across the four experiments ($F_{3,27} = 10.6$, $P < 0.0001$; Fig. 2a). *Post hoc* contrasts demonstrated that this interaction was partially derived from reliable increases in reaction time for the conflict conditions (mid and high) of response to feature to dimension to context tasks ($F > 6.4$, $P < 0.05$). Differences in error rates between patients and controls also followed an increasing pattern ($F_{3,27} = 5.2$, $P < 0.05$; see Supplementary Fig. 1 online). However, because sources of error were more variable, our analysis focused on correct trial reaction time.

The increasing difference in reaction time between patients and controls across tasks could reflect control deficits in two ways: higher-order control demands could increasingly challenge all patients, regardless of the site of their lesion, and so their performance could become differentially impaired as the task complexity increases, or deficits in higher-level tasks will be more likely across patients, regardless of the site of their lesion, than deficits at lower level tasks because of the asymmetric dependencies predicted by hierarchy, and so the larger deficits would reflect this aggregate likelihood. In the latter case, then the presence of an impairment at any level should increase the likelihood of an impairment at all higher levels, but should not increase the odds of an impairment at a lower level. This can be expressed as the change over the prior probability of a deficit at any level of the hierarchy, $p(D)$, when the probability of a deficit is conditioned on a deficit at any lower level, $p(D|L)$, versus a deficit at any higher level, $p(D|H)$.

A deficit was assigned for a task if a patient's average performance on mid/high conflict conditions was at least 2 s.d. worse than that of age-matched controls. The probability of a deficit on any task, $p(D)$, was 62% across the patients. Although there was a quantitative increase in the frequency of deficits as tasks required higher levels of control (Fig. 2b), these deficit frequencies were not reliably different across tasks ($F = 1.2$). Notably, however, the probability of a deficit at any level given a deficit at a lower level, $p(D|L)$, was 91% across patients, which was significantly different from $p(D)$ (Bayes factor (posterior odds/prior odds) = 6.7). In contrast, the probability of a deficit at any level given a deficit at a higher level, $p(D|H)$, was only 76%, which was a weak change over the prior probability (Bayes factor = 2.0). Notably, these results were not dependent on the 2 s.d. criterion for a deficit (see Supplementary Table 1 online). This asymmetry provides initial support for hierarchical dependencies among deficits at the different levels and the aggregation account of the group data.

Next, we considered whether these hierarchical deficit dependencies were related to the locus of damage along the rostro-caudal axis of the frontal lobes. An observer-independent method assigned patients to lesion overlap groups on the basis of their behavioral performance across the four tasks. Vectors were created that corresponded to the idealized behavior of a patient with a selective deficit at a particular hierarchical level, such as response, feature, dimension or context (Fig. 3a). These vectors served as regressors in a

multiple regression on each patient's reaction time differences from age-matched controls across all conditions of all experiments. The assumption of this multiple regression approach is that if a patient has damage encompassing more than one level of control, then their behavioral profile will be consistent with the linear sum of these two deficit profiles. When the resulting partial correlation coefficient associated with a particular regressor was positive and significant, favoring inclusion rather than exclusion ($P < 0.1$), a patient was assigned to that lesion overlap group.

The model assigned all but two of the lesion patients to the feature or dimension groups and one patient was assigned to both. The resulting lesion overlap maps clearly delineated adjacent, but separate, foci of maximal overlap along the rostro-caudal axis of the PFC for the feature versus dimension groups (Fig. 3b). The more caudal and dorsal focus of lesion overlap in the feature group, approximately pre-PMd, corresponds closely to the site of activation associated with the parametric effect of feature conflict from fMRI of healthy participants²⁸ (Fig. 3b). The more rostral and ventral focus of lesion overlap in the dimension group, straddling the IFS, corresponds almost precisely to the site of activation associated with the parametric effect of dimension conflict from fMRI²⁸ (Fig. 3b). The high degree of correspondence between the fMRI and patient lesion overlap results provides strong convergent support for the participation of these regions in cognitive control at different levels of abstraction.

Consideration of the behavioral profiles of the feature and dimension groups indicated that there was a crossover interaction that distinguishes the behavioral profiles of these two groups of patients (Fig. 4a,b). The feature group was intact relative to controls through all lower-level control conditions and the low-competition condition of the feature task ($F_{1,2} = 4.8$, $P < 0.05$). A deficit was evident for the mid and high conditions of the feature task and all conditions of the dimension and context tasks (Fig. 4a). In contrast, the dimension group was intact through the low-conflict condition of the dimension task, and deficient for the mid and high conditions of the dimension task ($F_{1,6} = 49.9$, $P < 0.0001$) and all conditions of the context task (Fig. 4a). Excluding the one patient assigned to both groups, these distinct profiles produced a crossover interaction between the feature and dimension groups ($F_{1,5} = 15.3$, $P < 0.05$; Fig. 4b,c).

It is notable that the mid-competition condition of the dimension task for the feature group showed a smaller difference from controls than the low- or high-competition conditions as a result of chance variation ($F_{1,2} = 4.4$, $P = 0.19$). However, the crossover interaction is not simply an artifact of this aspect of the experimental design. When only data from the high-versus low-competition conditions across the feature and dimension tasks were included in the analysis, the crossover pattern was still evident and showed a strong trend ($F_{1,5} = 5.6$, $P = 0.06$). Hence, the interaction does not appear to be restricted to the mid-competition condition. Likewise, the crossover interaction does not arise from a floor effect that obscures the differences between the low-, mid- and high-competition conditions for the feature group. Error rates across these conditions for the feature group were 16% and reaction times across these conditions were 3,268 ms (range: 2,060–4,477 ms), well below the response deadline (15 s).

DISCUSSION

These results demonstrate that performance deficits across frontal patients grow progressively worse as contingencies are added to an action decision and cognitive control operates at higher orders of abstraction. However, rather than deriving from a uniform pattern of progressive deficit in each patient, this pattern is the result of an asymmetric effect of a control deficit at a given level on higher-level control tasks. Specifically, the site of

damage resulting from stroke along the rostro-caudal axis of the frontal lobes results in a deficit at a predictable level of abstraction and in tasks requiring higher levels of control, but leaves performance on tasks requiring only lower levels of control intact.

The crossover interaction in behavioral performance by the patients on the feature and dimension tasks demonstrates that, although the tasks themselves are not independent, the control processors involved at each level of the hierarchy are independent, consistent with their spatial segregation. Specifically, feature-deficit patients are impaired on the dimension task not because they have difficulty deciding which dimension is relevant to their match decision (a third-order choice), but rather as a result of the subsequent determination of a response on the basis of the match relationship between the items (second-order choice). We know this because when we subtracted an estimate of the patients' ability to make this response selection on the basis of a match decision (the low-conflict dimension condition), there was no difference from controls on this task. This subtraction only works if the dimension processor can resolve conflict at the dimension level independent of the state (damaged or healthy) of the lower-level feature processor. As such, the data are consistent with a central property of a hierarchy, namely that controllers at higher levels operate independently from the status of lower-level processors. Thus, the reason that the feature-deficit patients failed in the dimension task was because of a feature-level deficit. However, the reason that the dimension-deficit patients failed at the same dimension task was because of a dimension-level deficit; they had difficulty deciding which dimension was relevant to their match decision (a third-order choice). Conversely, to the extent that higher-order control is not required by a task, lower-order processors may operate independently of higher-order control, which was evident in the intact performance of dimension patients on the feature task. Thus, the dimension and feature processors are independent, although the feature and dimension tasks themselves are not.

More broadly, the crossover interaction also provides rare direct lesion evidence for the widely assumed heterogeneity of function in the frontal lobes. For example, the crossover interaction may be consistent with distinct functions that have previously been associated with these regions across separate experiments related to conditional selection of stimulus information in pre-PMd versus selection of more abstract categorical information and high-level monitoring in IFS/mid-dorsolateral PFC⁴²⁻⁴⁷.

It is important to clarify that general difficulty, as in difficulty arising from any cause, be it increased abstraction or another factor that makes response times longer, cannot account for the results. Difficulty was manipulated in two ways in this study, both of which result in reaction time increases, but only one of which is related to the locus of damage along the rostro-caudal axis. First, there was the level of abstraction across the experiments, as additional contingencies were added to the action decision. Second, there was the degree of competition at a given level of contingency, which increased parametrically over three levels. Both of these factors produced increases in reaction time. However, the degree of abstraction was associated with the probability effects and the regional differences in lesion overlap. This is similar to the results from our previous fMRI experiment using this same task structure²⁸ that showed that the locus of activation along the rostro-caudal axis was related to the degree of abstraction and not to the level of competition. Instead, the level of competition selectively increased the amplitude of the fMRI response in a given region, but did not determine which region was activated along the rostro-caudal axis. Therefore, given that we manipulated two types of difficulty and located these effects related to only one, a single construct of difficulty cannot fully account for these or our previous results.

It should also be clarified that the lesion overlap approach taken here differs from other published approaches that define patient groups on the basis of their lesion overlap and then

assess any behavioral differences^{48,49}. Here, we defined our patient groups on the basis of a behavioral deficit profile and then looked at the regions of overlap among patients with a common profile. Although it is sometimes difficult to predict the behavioral profile of a particular patient on the basis of the location of their lesion, our results demonstrate that if a patient has a particular behavioral profile, then there is some consistency regarding the rostro-caudal locus of their lesion and how that patient will perform on other tasks requiring higher or lower levels of cognitive control.

Finally, our design does not address whether lower-order control processors are differentially affected by impairments in higher-order control when between-level interactions are required to complete a task. For example, higher-order decisions could modulate the degree of competition present at lower levels, as in constraining the number of lower-level choices by choosing a higher-level path. Such a test will be required to demonstrate that higher levels modulate lower levels, an important prediction in a strong processing hierarchy. Moreover, there may be feedback influences of lower-level on higher-level control. Our findings suggest that hierarchy may be a fruitful framework in which to understand frontal lobe architecture and systems-level processing and motivate further study.

METHODS

Patients and controls

We recruited 11 patients (average 56.6 years, range 45–73; 4 female) from the Northern California Veterans Administration Health Care System. An additional patient was recruited, but was unable to perform any of the tasks. Damage in all of the patients was caused by cerebral infarction of the middle cerebral artery. Testing took place at least 6 months post-stroke.

The extent of damage was assessed from structural MRI or computed tomography scans. Estimates were reconstructed on normalized templates by an expert anatomist who was blind to patient performance and then digitized to assess overlap. For visualization, digital masks were overlaid on a high-resolution MNI canonical image using MRICro (<http://www.sph.sc.edu/comd/rorden/mricro.html>).

We enrolled 43 control participants (26 female) following screening for any history of neurological or psychiatric disorder. The entire cohort of controls ranged in age from 21 to 69 years. For each patient, controls were selected from the cohort whose age was within 5 years of the lesion patient. From this selection procedure, a subset of 24 controls (12 female, ages 41–69) was included in the analysis. Patients and controls were thus matched for age (average difference was -0.14 years, $t_{10} = 0.34$) and years of education (average difference was 2.1 years, $t_{10} = 2.1$).

Patients and controls had normal or corrected-to-normal vision, and normal color vision, as verified by the Ishihara test for color deficiency. Informed consent was obtained in accordance with procedures approved by the Committees for Protection of Human Subjects at the University of California, Berkeley and the Northern California Veterans Administration Healthcare System. Participants were paid for their participation in the study.

Behavioral tasks

Patients and controls were tested on a battery of four response-selection tasks that were designed to test progressively higher degrees of hierarchically ordered control (Fig. 1a). These tasks were adapted from our previous fMRI experiment²⁸.

In the response task, participants viewed a series of colored boxes that were presented one at a time and selected a response on a keypad on the basis of the box color. Competition increased with the number of alternative responses on a given block of trials increasing from one (low) to two (mid) to four (high).

In the feature task, the series of colored boxes each contained a single object that varied from trial to trial along one perceptual dimension (either texture or orientation between subjects). The participants were required to decide whether a particular target feature along that dimension (that is, a rough texture) was presented on each trial. The participant made a positive response on the keypad if the target feature was present and a negative response to any other feature. The target feature that cued a positive response for a given trial was itself cued by the color of the box. Competition increased with the number of alternative target features for a given block of trials increasing from one (low) to two (mid) to four (high).

In the dimension task, the series of colored boxes each contained two objects that each varied along four dimensions (texture, shape, size and orientation) from trial to trial. The participants were required to decide whether the objects matched along only one of those dimensions on each trial. The relevant dimension was cued by the colored box. Competition increased with the number of alternative dimensions for a given block increasing from one (low) to two (mid) to four (high).

The context task was identical in terms of the task instructions to the dimension experiment, except that two dimensions were always relevant across all blocked conditions. Moreover, in the context task, a given color cue could map to different dimensions on different blocks (in the dimension task, a given color always mapped to one dimension). Thus, in the context task, it is necessary to use information about the current temporal frame (the current block or the most recent instructions) to select the appropriate mapping for a given color cue. Thus, competition was manipulated by varying the frequency across blocks that a given color cue (the context) mapped to a specific dimension. Certain color-to-dimension mappings were relevant for 100% of the blocks in which that cue was encountered, other color-to-dimension mappings were relevant for 50% of blocks in which that color was encountered and other color-to-dimension mappings were relevant on only 25% of blocks in which that color is encountered. In the latter two cases, determining which color-to-dimension mapping is currently relevant required the selection of a particular color-to-dimension mapping on the basis of the instructions of the current block. In this way, as the frequency of a given color-to-dimension mapping decreases, uncertainty or competition with other mappings increases and so selection of the currently relevant mapping requires more control.

Design and experimental parameters

The four tasks were tested across 2–4 sessions for each participant. To control for mapping frequencies, the context task was always performed first. The remaining three tasks were counterbalanced for order across participants with the constraint that at least one task (either feature or response) come before the performance of the dimension task.

The response, dimension and context tasks included 192 trials and the feature task included 186 trials. The response, feature and dimension tasks consisted of six blocks, two of each competition condition (low, mid and high), counterbalanced for order across participants. In the context task, there were 12 short blocks that permitted manipulation of mapping frequencies from low to mid to high across blocks. The order of blocks, cycled twice, was two blocks of low-competition conditions, followed by two blocks of high-competition conditions, and finally by two blocks of mid-competition conditions. This fixed order was provided so that participants could take advantage of order as an additional cue for selecting

the appropriate color-to-dimension mapping. The low, high and mid order was used to decouple fatigue or practice effects from the parametric manipulation of competition.

Individual trials in all experiments were self-paced up to a limit of 15 s. However, all participants were encouraged to respond as quickly and as accurately as possible on every trial. The specific color mappings, responses and objects used in the tasks were counterbalanced across subjects and two color sets were used to minimize confusion between tasks. Where applicable in each experiment, color cue, response, feature and dimension switches were controlled for frequency across blocks of each condition. All combinations of colors and features in the feature experiment and colors and shapes in the dimension and context experiments were controlled across competition and switching conditions.

Prior to performing each task, patients and controls were shown all the color mappings that they would encounter for that task, one block at a time. The mappings were covered and the participants were quizzed verbally. They then performed two practice blocks with the mapping set that they had just memorized. In the first practice block, the relevant mappings were available at the top of the screen, if a reminder was needed. The second practice block was identical to the experimental setting.

Data analysis

Median reaction time was obtained for each participant from correct trials. In cases in which a control subject performed greater than 2 s.d. above the mean reaction time of the entire control cohort ($n = 43$) or if their error rates were at chance for a given condition, that participant was excluded from the group average for that particular task.

We conducted a deficit probability analysis to determine the change over the prior probability of a deficit at any level, $p(D)$, given a deficit at a lower, $p(D|L)$, and at a higher, $p(D|H)$, level. First, the patient and control data were linearly corrected for simple motor speed by subtracting the reaction time for the response task low-competition condition from all other reaction times. This provides a measure of simple reaction time and estimates speed in the absence of cognitive control. The average reaction times of the mid and high conditions of each task were calculated for the patients and then standardized to a Z score on the basis of the matched control distribution. For standardized scores greater than 2 s.d., the corresponding task was coded as being deficient. For example, if the average reaction time for the mid and high conditions of the dimension task was 2 s.d. above the mean of controls for a particular patient, then this patient was listed as having a deficit at the dimension level. These deficit counts were then used to calculate the following probabilities: the probability of a deficit at each level, $p(\text{Response})$, $p(\text{Feature})$, $p(\text{Dimension})$ and $p(\text{Context})$, the base rate probability of a deficit at any level, $p(D)$, the conditional probability of a deficit at any level given a deficit at a lower level, $p(D|L)$, and the conditional probability of a deficit at any level given a deficit at a higher level, $p(D|H)$. The Bayes factor is the ratio of the posterior odds to the prior odds. The convention is that a Bayes factor less than 3 is considered to be negligible, a factor between 3 and 10 is substantial or implies supportive evidence, and a factor above 10 is considered to be strong evidence⁵⁰. This analysis was also conducted for deficit criteria ranging from 1–2.5 s.d. (see Supplementary Table 1).

We used an observer-independent overlap method to assign patients to lesion overlap groups on the basis of their behavioral performance across the tasks. Regressors were created that reflected the predicted deficits for patients in each of the four groups, response, feature, dimension and context (Fig. 3a). These predictions derived from the hierarchy hypothesis and made three assumptions: conditions that include competition only at levels below the level of deficit will be intact, performance at the level of deficit will be worse depending on

the degree of competition at that level (thus, performance will get parametrically worse with parametric increases in competition) and performance on conditions that include competition at higher levels will also be impaired.

These prediction vectors were then included in a multiple regression. Patients who had a reliable and positive partial correlation coefficient were included in a particular overlap group. A lenient threshold was used for inclusion ($P < 0.1$) to include as many patients as possible in the overlap maps. Overlap masks were created using MRICro on the basis of the normalized lesion masks generated for each patient. The mask of one patient with a right-sided lesion was mirror flipped to permit its comparison with the left sided lesions of the other patients in the group. Behavioral averages from the group assignments, after excluding the one subject included in both groups, were used to compute the behavioral crossover interaction.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We are grateful to R.T. Knight and D. Scabini for their help with patient recruitment and lesion characterization. We also would like to thank W. Heindel and A. Kayser for their input on revisions of this manuscript. This work was supported by the US National Institutes of Health (grants MH63901 and NS40813), the Veterans Administration Research Service and a National Research Service Award (F32 NS053337).

References

1. Badre D, Wagner AD. Selection, integration and conflict monitoring; assessing the nature and generality of prefrontal cognitive control mechanisms. *Neuron* 2004;41:473–487. [PubMed: 14766185]
2. D'Esposito M, et al. The neural basis of the central executive system of working memory. *Nature* 1995;378:279–281. [PubMed: 7477346]
3. D'Esposito M, Postle BR, Rypma B. Prefrontal cortical contributions to working memory: evidence from event-related fMRI studies. *Exp Brain Res* 2000;133:3–11. [PubMed: 10933205]
4. Duncan J. An adaptive coding model of neural function in prefrontal cortex. *Nat Rev Neurosci* 2001;2:820–829. [PubMed: 11715058]
5. Fuster, JM. *The Prefrontal Cortex: Anatomy, Physiology and Neuropsychology of the Frontal Lobe*. Lippincott-Raven Publishers; Philadelphia, Pennsylvania: 1997.
6. Goldman-Rakic PS. The prefrontal landscape: implications of functional architecture for understanding human mentation and the central executive. *Phil Trans R Soc Lond B* 1996;351:1445–1453. [PubMed: 8941956]
7. Koechlin E, Summerfield C. An information theoretical approach to prefrontal executive function. *Trends Cogn Sci* 2007;11:229–235. [PubMed: 17475536]
8. Miller EK, Cohen JD. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 2001;24:167–202. [PubMed: 11283309]
9. Passingham, RE. *The Frontal Lobes and Voluntary Action*. Oxford University Press; Oxford: 1993.
10. Passingham, RE.; Rowe, JB. Dorsal prefrontal cortex: maintenance in memory or attentional selection?. In: Stuss, DT.; Knight, RT., editors. *Principles of Frontal Lobe Function*. Oxford University Press; Oxford: 2002. p. 221-232.
11. Petrides M. Lateral prefrontal cortex: architectonic and functional organization. *Phil Trans R Soc Lond B* 2005;360:781–795. [PubMed: 15937012]
12. Shallice T, Burgess PW. Deficits in strategy application following frontal lobe damage in man. *Brain* 1991;114:727–741. [PubMed: 2043945]

13. Smith EE, Jonides J. Storage and executive processes in the frontal lobes. *Science* 1999;283:1657–1661. [PubMed: 10073923]
14. Stuss, DT.; Benson, DF. The frontal lobes and control of cognition and memory. In: Percecman, E., editor. *The Frontal Lobes Revisited*. The IRBN Press; New York: 1987. p. 141-158.
15. Bunge SA. How we use rules to select actions: a review of evidence from cognitive neuroscience. *Cogn Affect Behav Neurosci* 2004;4:564–579. [PubMed: 15849898]
16. Carter CS, et al. Anterior cingulate cortex, error detection and the on-line monitoring of performance. *Science* 1998;280:747–749. [PubMed: 9563953]
17. O'Reilly RC, Frank MJ. Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. *Neural Comput* 2006;18:283–328. [PubMed: 16378516]
18. Courtney, SM.; Roth, JK.; Sala, JB. A hierarchical biased-competition model of domain-dependent working memory maintenance and executive control. In: Osaka, N.; Logie, R.; D'Esposito, M., editors. *The Cognitive Neuroscience of Working Memory*. Oxford University Press; Oxford: 2007. p. 369-383.
19. Cohen JD, Dunbar K, McClelland JL. On the control of automatic processes: a parallel distributed processing account of the Stroop effect. *Psychol Rev* 1990;97:332–361. [PubMed: 2200075]
20. Cohen JD, Servan-Schreiber D. Context, cortex and dopamine: a connectionist approach to behavior and biology in schizophrenia. *Psychol Rev* 1992;99:45–77. [PubMed: 1546118]
21. Desimone R, Duncan J. Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 1995;18:193–222. [PubMed: 7605061]
22. MacDonald AW III, Cohen JD, Stenger VA, Carter CS. Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science* 2000;288:1835–1838. [PubMed: 10846167]
23. Wallis JD, Anderson KC, Miller EK. Single neurons in prefrontal cortex encode abstract rules. *Nature* 2001;411:953–956. [PubMed: 11418860]
24. Botvinick MM. Multilevel structure in behaviour and in the brain: a model of Fuster's hierarchy. *Philos Trans R Soc Lond B Biol Sci* 2007;362:1615–1626. [PubMed: 17428777]
25. Botvinick MM. Hierarchical models of behavior and prefrontal function. *Trends Cogn Sci* 2008;12:201–208. [PubMed: 18420448]
26. O'Reilly RC, Noelle DC, Braver TS, Cohen JD. Prefrontal cortex and dynamic categorization tasks: representational organization and neuromodulatory control. *Cereb Cortex* 2002;12:246–257. [PubMed: 11839599]
27. Hazy TE, Frank MJ, O'Reilly RC. Towards an executive without a homunculus: computational models of the prefrontal cortex/basal ganglia system. *Philos Trans R Soc Lond B Biol Sci* 2007;362:1601–1613. [PubMed: 17428778]
28. Badre D, D'Esposito M. Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. *J Cogn Neurosci* 2007;19:2082–2099. [PubMed: 17892391]
29. Christoff, K.; Keramatian, K. Abstraction of mental representations: theoretical considerations and neuroscientific evidence. In: Bunge, SA.; Wallis, JD., editors. *Perspectives on Rule-Guided Behavior*. Oxford University Press; New York: 2007.
30. Koechlin E, Jubault T. Broca's area and the hierarchical organization of human behavior. *Neuron* 2006;50:963–974. [PubMed: 16772176]
31. Koechlin E, Ody C, Kouneiher F. The architecture of cognitive control in the human prefrontal cortex. *Science* 2003;302:1181–1185. [PubMed: 14615530]
32. Badre D. Cognitive control, hierarchy and the rostro-caudal organization of the frontal lobes. *Trends Cogn Sci* 2008;12:193–200. [PubMed: 18403252]
33. Buckner RL. Functional-anatomic correlates of control processes in memory. *J Neurosci* 2003;23:3999–4004. [PubMed: 12764084]
34. Fuster JM. The prefrontal cortex—an update: time is of the essence. *Neuron* 2001;30:319–333. [PubMed: 11394996]
35. Rowe JB, et al. Is the prefrontal cortex necessary for establishing cognitive sets? *J Neurosci* 2007;27:13303–13310. [PubMed: 18045924]

36. Sakai K, Passingham RE. Prefrontal interactions reflect future task operations. *Nat Neurosci* 2003;6:75–81. [PubMed: 12469132]
37. Sakai K, Passingham RE. Prefrontal set activity predicts rule-specific neural processing during subsequent cognitive performance. *J Neurosci* 2006;26:1211–1218. [PubMed: 16436608]
38. Barbas H, Pandya DN. Architecture and frontal cortical connections of the premotor cortex (area 6) in the rhesus monkey. *J Comp Neurol* 1987;256:211–228. [PubMed: 3558879]
39. Petrides M, Pandya DN. Efferent association pathways from the rostral prefrontal cortex in the macaque monkey. *J Neurosci* 2007;27:11573–11586. [PubMed: 17959800]
40. Christoff K, Gabrieli JDE. The frontopolar cortex and human cognition: evidence for a rostrocaudal hierarchical organization within the human prefrontal cortex. *Psychobiology* 2000;28:168–186.
41. Christoff K, Ream JM, Geddes LPT, Gabrieli JDE. Evaluating self-generated information: anterior prefrontal contributions to human cognition. *Behav Neurosci* 2003;117:1161–1168. [PubMed: 14674837]
42. Halsband U, Passingham RE. Premotor cortex and the conditions for movement in monkeys (*Macaca fascicularis*). *Behav Brain Res* 1985;18:269–277. [PubMed: 4091963]
43. Hampshire A, Duncan J, Owen AM. Selective tuning of the blood oxygenation level-dependent response during simple target detection dissociates human frontoparietal subregions. *J Neurosci* 2007;27:6219–6223. [PubMed: 17553994]
44. Hampshire A, Thompson R, Duncan J, Owen AM. The target selective neural response—similarity, ambiguity, and learning effects. *PLoS ONE* 2008;3:e2520. [PubMed: 18575585]
45. Heekeren HR, Marrett S, Bandettini PA, Ungerleider LG. A general mechanism for perceptual decision-making in the human brain. *Nature* 2004;431:859–862. [PubMed: 15483614]
46. Petrides M, Alivisatos B, Evans AC, Meyer E. Dissociation of human mid-dorsolateral from posterior dorsolateral frontal cortex in memory processing. *Proc Natl Acad Sci USA* 1993;90:873–877. [PubMed: 8430100]
47. Rushworth MF, et al. Attentional selection and action selection in the ventral and orbital prefrontal cortex. *J Neurosci* 2005;25:11628–11636. [PubMed: 16354921]
48. Dreher JC, Koechlin E, Tierney M, Grafman J. Damage to the fronto-polar cortex is associated with impaired multitasking. *PLoS ONE* 2008;3:e3227. [PubMed: 18795100]
49. Thompson-Schill SL, et al. Verb generation in patients with focal frontal lesions: a neuropsychological test of neuroimaging findings. *Proc Natl Acad Sci USA* 1998;95:15855–15860. [PubMed: 9861060]
50. Jeffreys, H. *Theory of Probability*. Clarendon Press; Oxford: 1961.

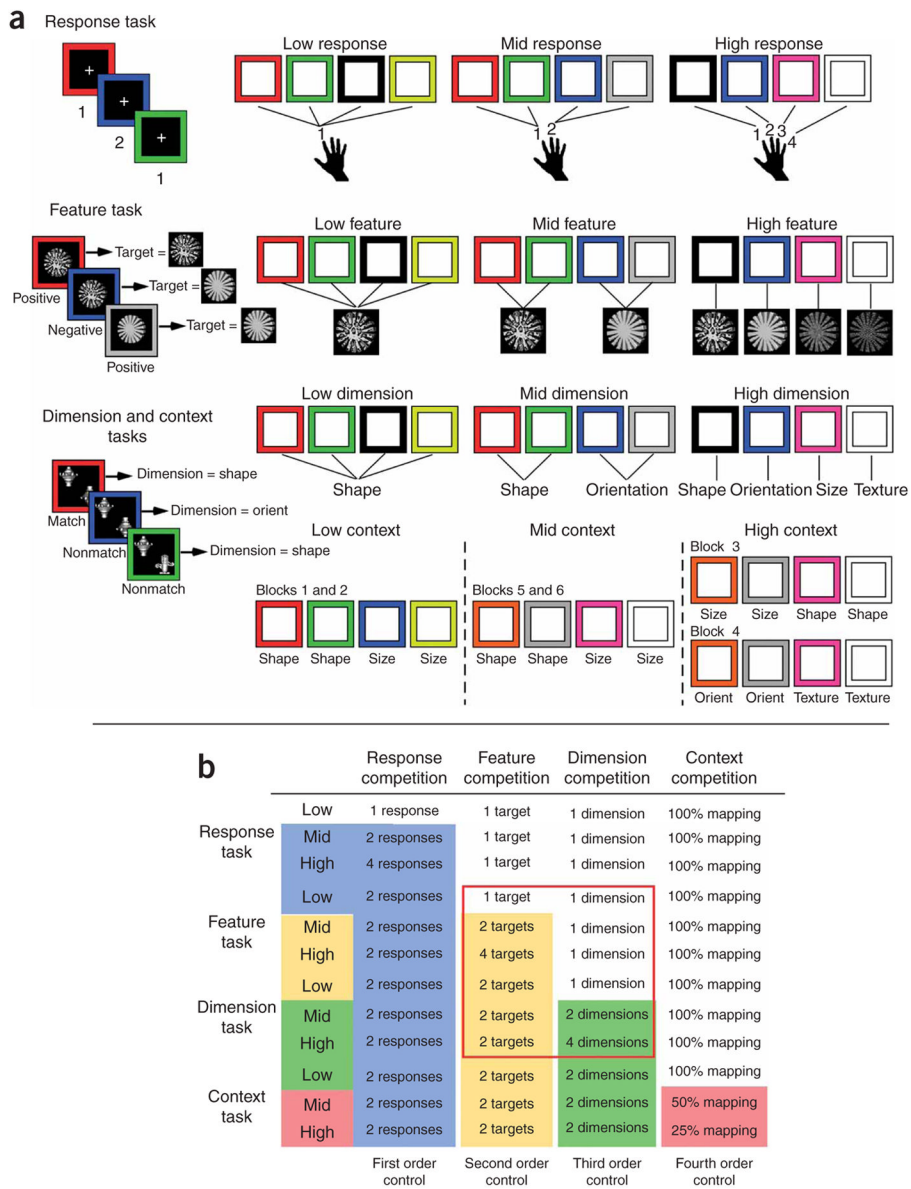
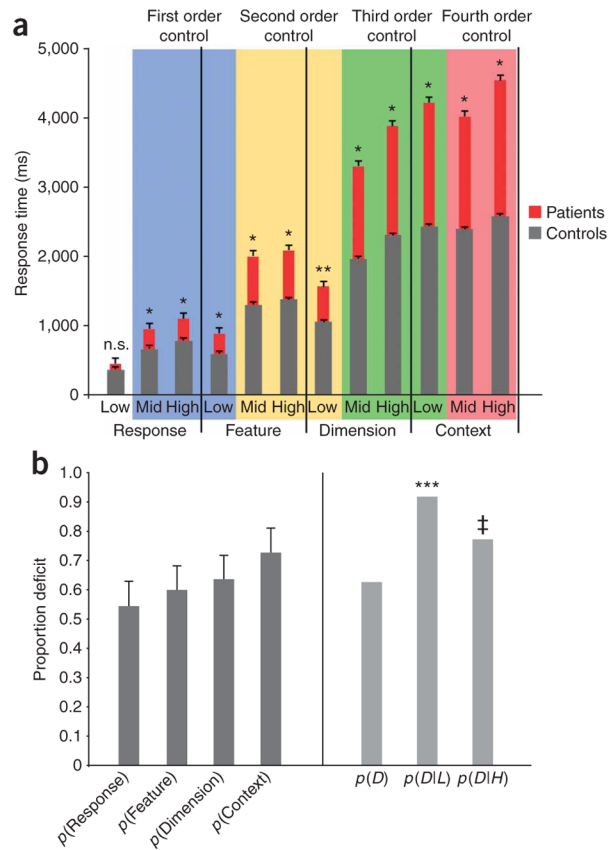


Figure 1. Trial events and task analysis of the four response-selection tasks. **(a)** On each trial of the response task, participants chose a response key on the basis of the color of a presented square. Competition conditions were low (one response), mid (two alternative responses) and high (four alternative responses). On each trial of the feature task, the participant looked for a particular target feature (for example, a mottled texture) based on the color of the square. They made a positive response if the target feature was presented and a negative response otherwise. Competition conditions included one target feature (low), two alternative target features (mid) or four alternative target features (high). Logically, this manipulation increases the number of sets of response mappings from one to four. Thus, the number of targets may be thought of as the number of response sets. On each trial of the dimension and context tasks, the participant decided whether two objects matched along a particular dimension (for example, shape) that was cued by the color of the square. Dimension competition conditions were one dimension (low), two alternative dimensions

(mid) or four alternative dimensions (high). During the context experiment, there were always two alternative dimensions, but competition was introduced by decreasing the frequency with which a given color mapped to a given dimension (low, 100%; mid, 50%; high, 25% mapping frequency). Thus, by definition, from the first order through the fourth order of the hierarchy, competition was defined by the number of responses, targets, dimensions and mappings, respectively. **(b)** A task analysis depicts the nested hierarchical relationship in control demands (columns) among the four tasks (rows). Color-coding highlights conditions for which competition at the response (blue), feature (yellow), dimension (green) or context (red) levels was present. Thus, this table indicates how control demands at different levels accumulate as each level of contingency is added in each task. Also, note that the low-competition condition of each task is equivalent in control demands to the mid condition of the task one level subordinate. Finally, the red outline highlights the conditions permitting a crossover interaction.

**Figure 2.**

Overall performance across the four tasks. **(a)** Reaction time for patients (red) and controls (gray) is plotted across the four tasks that increase, from left to right, in degree (low, mid and high) and order of competition (response, feature, dimension and context). Color-coding indicates conditions across tasks that include equivalent levels of conflict at the response (blue), feature (yellow), dimension (green) or context (red) levels but no higher-order competition. Notably, the difference between patients and controls grew as higher-order control was required ($*P < 0.05$, $**P = 0.06$, error bars represent s.e.m.). **(b)** The proportion of patients showing deficits in each task grew quantitatively, but not reliably, from response to context (left). Notably, however, the probability of a deficit at any level, $p(D)$, was reliably greater when conditioned on a deficit at any lower level, $p(D|L)$, relative to when it was conditioned on a deficit at any higher level, $p(D|H)$. (***) indicates Bayes factor of 6.7, ‡ indicates Bayes factor of 2, error bars represent s.e.m.)

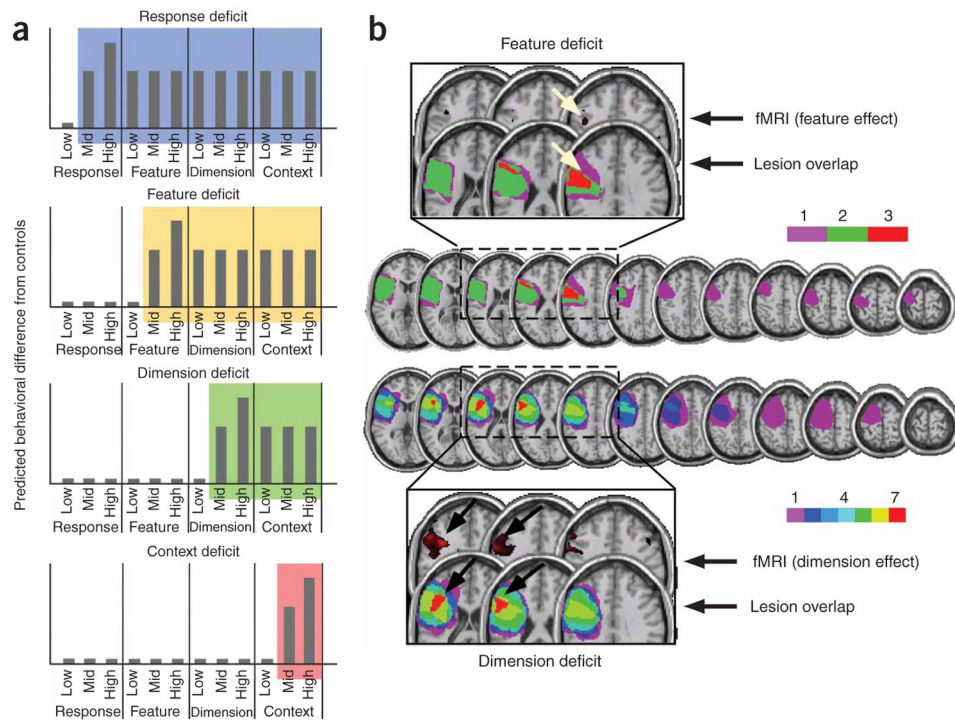


Figure 3. Observer-independent overlap analysis. **(a)** Regressors were generated representing the pattern of data across tasks and conditions for an idealized deficit at each level of control on the basis of the asymmetrical hierarchical assumptions. Bars indicate difference from controls in arbitrary units. Color-coding highlights the conditions for which deficits should emerge for patients with impairments in response (blue), feature (yellow), dimension (green) and context (red) control. **(b)** Results from the lesions overlap analysis revealed a distinction in the peak of overlap (red) among dimension patients around the IFS/dorsolateral prefrontal cortex and the peak of overlap (red) among feature patients in anterior dorsal premotor cortex. Color bar indicates the number of patients contributing to each colored region. Insets show correspondence between sites of lesion overlap from this study and the activation associated with the parametric effect of dimension (top) and feature (bottom) conflict from ref. ²⁸. Arrows on slices are in the same position for precise comparison.

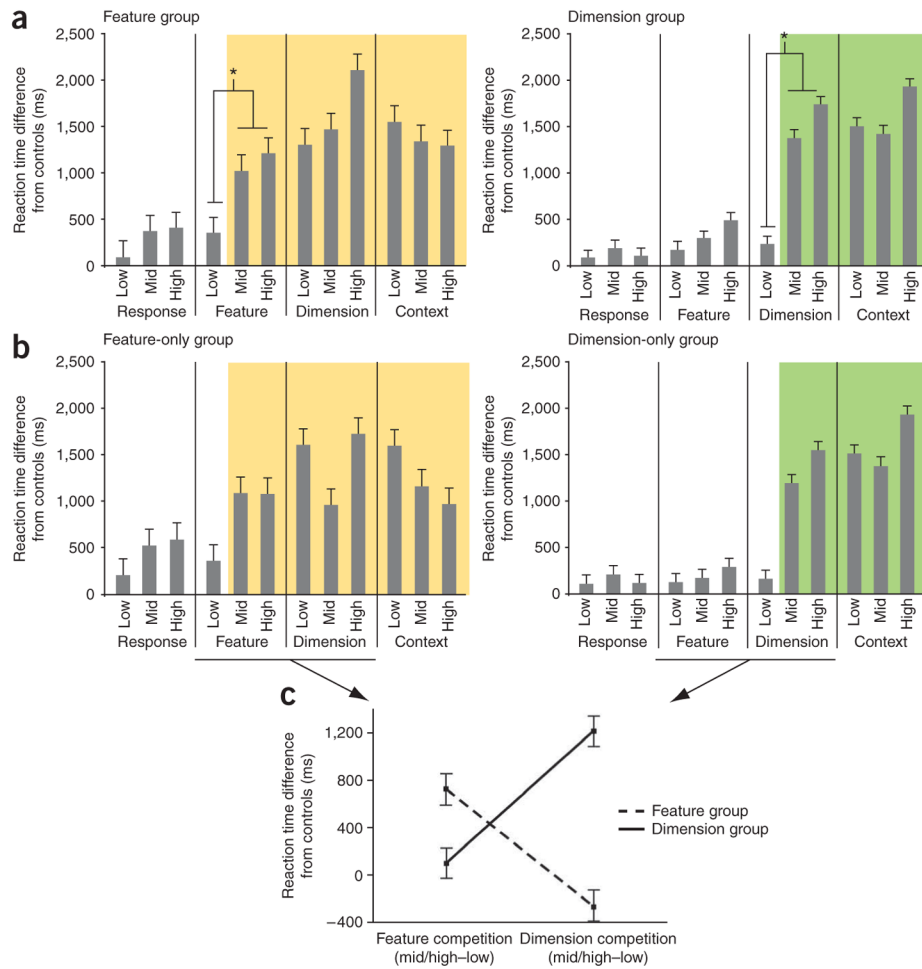


Figure 4. Performance of dimension and feature patient groups. **(a)** The differences in reaction time between patients and controls in the feature (left) and dimension (right) groups are plotted across competition conditions and tasks. Colored shading highlights occurrences of a reliable stepwise increase in a feature (yellow) or dimension (green) control deficit ($*P < 0.05$). **(b)** The differences in reaction time between patients and controls in the feature (left) and dimension (right) groups, excluding the one patient that was categorized as having both feature and dimension deficits, are plotted across competition conditions and tasks. **(c)** The differences from controls in the reaction time change between conflict (mid/high) and nonconflict (low) conditions of the feature (left) and dimension (right) tasks are plotted for the feature-only (blue) and dimension-only (red) overlap groups. The crossover interaction supports a double dissociation between these groups. All error bars represent s.e.m.