Using Neural Pattern Classifiers to Quantify the Modularity of Conflict–Control Mechanisms in the Human Brain

Jiefeng Jiang[1,2] and Tobias Egner[1,2]

[1]Center for Cognitive Neuroscience and [2]Department of Psychology and Neuroscience, Duke University, Durham, NC, USA

Address correspondence to Tobias Egner, Center for Cognitive Neuroscience, Duke University, LSRC Box 90999, Durham, NC 27708, USA.
Email: tobias.egner@duke.edu

Resolving conflicting sensory and motor representations is a core function of cognitive control, but it remains uncertain to what degree control over different sources of conflict is implemented by shared (domain general) or distinct (domain specific) neural resources. Behavioral data suggest conflict–control to be domain specific, but results from neuroimaging studies have been ambivalent. Here, we employed multivoxel pattern analyses that can decode a brain region's informational content, allowing us to distinguish incidental activation overlap from actual shared information processing. We trained independent sets of "searchlight" classifiers on functional magnetic resonance imaging data to decode control processes associated with stimulus-conflict (Stroop task) and ideomotor-conflict (Simon task). Quantifying the proportion of domain-specific searchlights (capable of decoding only one type of conflict) and domain-general searchlights (capable of decoding both conflict types) in each subject, we found both domain-specific and domain-general searchlights, though the former were more common. When mapping anatomical loci of these searchlights across subjects, neural substrates of stimulus- and ideomotor-specific conflict–control were found to be anatomically consistent across subjects, whereas the substrates of domain-general conflict–control were not. Overall, these findings suggest a hybrid neural architecture of conflict–control that entails both modular (domain specific) and global (domain general) components.

Keywords: cognitive control, conflict, multivoxel pattern analysis

Introduction

A fundamental challenge to goal-directed behavior is that a multitude of stimuli vie for control over our actions. "Conflict–control", the ability to resolve this competition in line with an organism's current goals, is therefore considered a core executive function (Norman and Shallice 1986; Desimone and Duncan 1995; Botvinick et al. 2001; Miller and Cohen 2001). A key question in this regard is whether the brain is co-opting a singular, central resource for dealing with different types of conflicting representations (domain-general conflict–control), or whether we have evolved distinct adaptations to confront the specific challenges posed by competition in different processing pathways (domain-specific conflict–control; Egner 2008). For instance, would conflict that arises from activation of competing stimulus representations be resolved by the same resource as conflict that emanates from coactivation of competing response codes? Recent behavioral evidence indicates that conflict–control displays a modular, domain-specific organization (Wendt et al. 2006; Egner et al. 2007; Funes et al. 2010), with separate conflict–control mechanisms being responsive to distinct domains and capable of running in parallel with each other (Egner 2008). On the other hand, functional magnetic resonance imaging (fMRI) studies

manipulating either different conflict sources or components have produced ambivalent results, ranging from complete domain specificity (van Veen and Carter 2005; Liston et al. 2006; Egner et al. 2007) to a mix of domain-specific and domain-general activations (Milham et al. 2001; Fan et al. 2003; Liu et al. 2004; Kim et al. 2010, 2011) to completely overlapping activations across conflict domains (Peterson et al. 2002).

However, the degree to which regional activation overlap across conflict domains assayed in previous fMRI studies accurately describes shared conflict–control processes is uncertain. First, the amount of activation overlap across domains depends on arbitrary statistical thresholding. Secondly, it is unknown whether regions of shared activations carry essential information about the conflict–control process itself or serve incidental functions. Thirdly, common regional fMRI signal may reflect distinct contributions from interspersed, domain-specific neuronal populations. Fourthly, voxel ensembles that do not display mean activation differences between conditions may nevertheless code for essential information, such that key regions engaged in domain-general or -specific conflict–control could pass undetected in traditional (univariate) fMRI analyses. Conversely, higher mean activation in one condition compared with another does not necessarily translate into higher informational content (Kok et al. 2012).

In order to overcome these limitations, we turned to fMRI multivoxel pattern analysis (MVPA), which harnesses multivariate activity patterns over spatially distributed voxels to train classifiers decoding whether a particular brain region carries diagnostic information about specific stimuli or cognitive processes (Haxby et al. 2001; Haynes and Rees 2006; Kriegeskorte et al. 2006; Norman et al. 2006). MVPA operates in reference to a nonarbitrary threshold (chance-level classification); it increases effective spatial resolution, thus allowing for the detection of interspersed neuronal populations of varying functional characteristics (Kamitani and Tong 2005; Esterman et al. 2009); and it enables the detection of task-relevant population codes that are not characterized by mean activation differences between conditions (Howard et al. 2009).

We therefore combined whole-brain searchlight MVPA (Kriegeskorte et al. 2006; Haynes et al. 2007) with a task that independently varied "stimulus-conflict", the competition between task-relevant and -irrelevant perceptual or semantic stimulus features (Kornblum et al. 1990), and "ideomotor-conflict", the competition between a (task-irrelevant) motor affordance of a stimulus and the motor act that is required in response to that stimulus (Shin et al. 2010). Using independent classifiers trained on multivoxel signal patterns, we defined domain-general conflict–control regions as searchlights that are capable of decoding both stimulus- and ideomotor-conflicts, whereas domain-specific control–control

regions were defined as searchlights that could only successfully decode one type of conflict. Subsequently, we quantified the proportion of domain-general versus domain-specific processors in each subject and mapped the anatomical loci of these searchlights across subjects. The results produced evidence for both domain-specific and domain-general processing, though with a substantial preponderance of the former. Moreover, whereas anatomical locations of stimulus- and ideomotor-conflict–control searchlights were consistent across subjects, regions supporting domain-general conflict–control were detected in the group analysis only at more lenient (uncorrected) statistical thresholds.

Materials and Methods

Subjects

Twenty-nine healthy volunteers (mean age = 21.3, 15 females) gave informed consent in accordance with institutional guidelines. All subjects were native or highly proficient English speakers and had normal or corrected-to-normal vision. Five subjects were excluded due to low accuracy on the behavioral task (<80% correct), and a further 3 subjects due to excessive movement (>3 mm or >3°) during imaging. The final sample consisted of 11 females and 10 males.

Apparatus and Stimuli

Stimulus delivery and behavioral data collection were carried out using Presentation software (http://www.neurobs.com/). Stimuli were presented on a back projection screen viewed via a mirror attached to the headcoil of the MRI scanner, and responses were collected using an MRI-compatible button box. Stimuli consisted of a collection of 24 black and white photographs of male and female faces (12 each) of neutral expression that were overlaid with red gender word labels ("male" and "female"), which could be printed in lower or upper case lettering, resulting in a total of 96 unique pictorial stimuli (Fig. 1A).

Procedure

On each trial, one face-word compound stimulus (subtending approximately 3° of horizontal and 4° of vertical visual angle) was presented against a gray background in either the left upper, left lower, right upper, or right lower quadrant of the screen (centered at approximately 7° of horizontal eccentricity from the vertical midline, and at approximately 3° of vertical eccentricity from the horizontal midline of the display), with a fixation cross occupying the center of the screen (Fig. 1A). Stimuli were presented for 1 s, followed by a jittered interstimulus interval ranging from 3 to 4.5 s in uniformly distributed steps of 500 ms, during which the fixation cross remained on screen. Subjects performed a speeded button response that categorized the gender of the face stimulus with either index finger (for example, left-hand response to male faces, right-hand response to female faces, counterbalanced across subjects), while trying to ignore the task-irrelevant gender labels and stimulus locations. Following a brief practice run to ensure subjects comprehended the task requirements, the experimental trials were presented across 4 scan runs, each entailing 96 trials.

Following the main task, subjects also performed a localizer task to define the fusiform face area (FFA; Kanwisher et al. 1997), a visual brain regions specialized in the processing of face stimuli. In this localizer task, subjects responded to direct repetitions of visual stimuli (1-back task) by pressing a button using the right index finger. Face and house stimuli were presented separately in 5 blocks each. Each block contained 15 stimuli, including 1 or 2 direct repetitions. Stimuli were presented at the center of the screen (subtending approximately 3° of horizontal and 4° of vertical visual angles) against a gray background for 750 ms. Each stimulus was then followed by a fixation cross shown at the center of the screen for 250 ms. The blocks presenting face and house stimuli alternated and were further randomly interleaved with 5 fixation blocks of 15 s each.
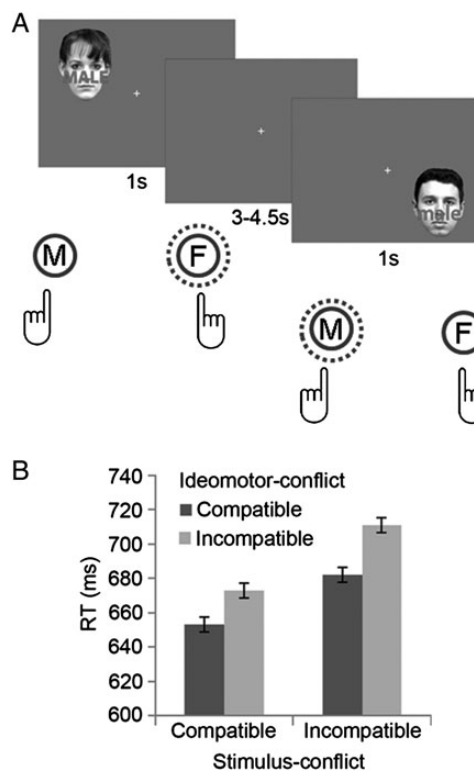
Figure 1. Experimental protocol and behavioral data. (A) Example stimuli and timing (and correct responses) of the experimental protocol. Assuming a button mapping of a left-hand response to male faces and a right-hand response to female faces (mappings were counterbalanced across subjects), the example depicts one trial that was both stimulus- and ideomotor-incompatible, followed by another trial that was stimulus-compatible, but ideomotor-incompatible. (B) RT data (±mean standard error [MSE]), plotted as a function of stimulus- and ideomotor-compatibility, exhibit additive main effects of these 2 factors.

Experimental Design

The experimental design of the main task integrates variants of the classic Stroop task (Stroop 1935) and Simon task (Simon 1969), by combining a task-relevant stimulus dimension (face gender) with 2 task-irrelevant stimulus dimensions that could give rise to stimulus- or ideomotor-conflict, respectively. Specifically, the semantic meaning of the gender label written across the face stimuli could be either compatible (50% of trials) or incompatible with the gender of the face stimulus, representing a potential source of stimulus-conflict (Kornblum et al. 1990; Egner et al. 2008). Moreover, the horizontal stimulus location (left/right) could be either compatible (50% of trials) or incompatible with the horizontal position of the correct response effecter, representing a potential source of ideomotor-conflict (Kornblum et al. 1990; King et al. 2010). Stimulus- and ideomotor-compatibility were varied independently of each other, resulting in a fully balanced 2 × 2 (stimulus domain: Compatible vs. incompatible × ideomotor domain: Compatible vs. incompatible) within-subject factorial design. To control for possible associative influences on behavior and imaging data stemming from stimulus feature repetitions across sequential trials, all stimulus features changed from each trial to the next, while the response feature (left or right button press) was equally likely to repeat or change from trial to trial. Specifically, no particular face stimulus was ever repeated over 2 successive trials, the gender labels alternated between lower case and upper case lettering across trials, and the spatial quadrant of stimulus presentation also always varied from trial to trial.

Behavioral Data Analysis

The behavioral signature of conflict–control we employed is given by the size of the interference (or conflict) effect, the differential in

response time (RT) to incompatible when compared with compatible stimuli. This interference score captures both the generation/detection and the within-trial resolution of conflict (given that we only analyzed trials where a correct answer was given). Note that an approximate segregation of conflict detection versus conflict resolution processes could in principle be achieved by analyzing first-order congruency sequence (conflict adaptation) effects (Gratton et al. 1992; Botvinick et al. 2001; Egner 2007). However, since this analysis results in a substantial reduction in trial count (going from 4 to 16 cells in the factorial design), which would have greatly reduced the power of the neural pattern classification, we opted to focus exclusively on the standard trial *n* effects. An added benefit of this approach is that it coincides with the way that the large majority of previous fMRI studies manipulating different conflict sources have analyzed their data, thus rendering the present MVPA results more comparable with prior, univariate fMRI results. The mean RT was computed in each subject for each of the experimental cells, excluding incorrect trials and RT values that deviated >2 standard deviations from an individual subject's grand mean. The trimmed RT values were then averaged across subjects and entered into repeated-measures 2-way analyses of variance (ANOVAs) with the factors stimulus-compatibility (compatible vs. incompatible) and ideomotor-compatibility (compatible vs. incompatible). Equivalent analyses were carried out on mean accuracy rates.

Image Acquisition
Images were acquired on a 3-T GE scanner (Milwaukee, WI, USA). Structural images were scanned using a T_1-weighted axial sequence parallel to anterior commissure–posterior commissure line (120 slices, slice thickness = 1 mm, time repetition [TR] = 8.124 ms, field of view (FoV) = 256 mm × 256 mm, in-plane resolution = 1 mm × 1 mm). Functional images were scanned using a T_2*-weighted inverse spiral sequence (40 slices, slice thickness = 3 mm, TR = 1.5 s, time echo [TE] = 24 ms, flip angle = 85°, FoV = 192 mm × 192 mm, in-plane resolution = 3 mm × 3 mm). The functional images of the main task were acquired in 4 runs of 330 images each. The functional images of the localizer task were acquired in one run of 314 images.

Image Preprocessing
Preprocessing of functional images was performed using SPM8 (http://www.fil.ion.ucl.ac.uk/spm/). For each run, the first 7 scans were discarded from further analysis. The remaining images were realigned to their mean image and corrected for differences in slice-time acquisition. Each subject's structural image was coregistered to the mean functional image and normalized to the Montreal Neurological Institute (MNI) template brain. The transformation parameters of the structural image normalization were then applied to the functional images. The normalized functional images retained their native spatial resolution.

General Linear Model Univariate Analysis
For the standard general linear model (GLM) analysis, normalized images were smoothed using an 8-mm Gaussian kernel. Note that the GLM results reported below were not strongly dependent on our choice of spatial smoothing kernel (8 mm³), as we obtained qualitatively equivalent results with a smaller kernel (6 mm³) and when using unsmoothed data. For each subject, a model of the main task was created via vectors of stimulus onsets corresponding to the 4 experimental trial types, along with vectors for error trials, head-motion parameters, and grand means of each run. Vectors were convolved with SPM8's canonical hemodynamic response function to produce a design matrix, against which the blood oxygenation level-dependent signal at each voxel was regressed. Within-subject effects of stimulus- (or ideomotor-) compatibility were assessed by contrasting activations between stimulus- (or idemotor-) incompatible and compatible regressors at each voxel. The resulting single-subjects contrast maps were then entered into group-level, 1-sample, *t*-tests, where subjects were treated as random effects. For the localizer task, boxcar vectors of onsets and durations were created for face- and house-blocks. These

vectors, along with 7 other vectors, modeling head-motion parameters and the grand mean were convolved with SPM8's canonical hemodynamic response function to produce the design matrix, which was then estimated at each voxel for each subject. Face-sensitive regions were defined through a face > house block contrast at the single-subject level. These contrast maps were then entered into group-level, 1-sample *t*-tests, where subjects were treated as random effects.

MVPA Feature Generation
The preprocessed fMRI images scanned during the main task also underwent an MVPA procedure (Fig. 2). For MVPA feature generation, we employed a strategy based on grouping a small number of trials within each condition in each run together in order to create "trial exemplars" (Misaki et al. 2010). Specifically, for each subject, the 96 trials of each experimental condition were randomly grouped into 24 exemplars of 3–4 trials (depending on exclusions due to error trials) per run, in order to reduce noise. There was no overlap of trials across exemplars, that is, each trial belonged to a single exemplar only. Feature extraction was performed via GLM analysis of unsmoothed data. The resulting design matrix contained 96 unique regressors (one for each "exemplar", of which there were 6 for each of the 4 stimulus- × ideomotor-conflict congruency conditions per each of 4 runs), 24 head-motion parameters (6 for each run) and 4 grand means (1 for each run). The *t*-value of each regressor was then estimated at each voxel and used as the feature of that regressor's corresponding exemplar. All features from the same run were normalized at each voxel to reduce variances of features between runs.

Searchlight MVPA
To decode stimulus- and ideomotor-compatibility using regional information content, a whole-brain (constrained by a gray matter mask derived from the segmentation of the SPM8 T1 template) searchlight MVPA (Kriegeskorte et al. 2006; Haynes et al. 2007) was conducted on the features. Similar to previous studies (Clithero et al. 2009), the searchlight was a spherical cluster with a radius of 3 voxels and contained up to 123 cortical voxels. Linear support vector machines (SVMs) were used. Specifically, we used Matlab functions "svmtrain" and "svmclassify" for training and testing an SVM, respectively. The box constrain-value controlling the soft margin was set to 1 (the default value) for all SVMs. The performance of SVMs was evaluated with a leave-one-run-out cross-validation approach. Four types of classifiers were trained on searchlights centered at each voxel, namely "pure ideomotor-conflict", "pure stimulus-conflict", "full ideomotor-conflict", and "full stimulus-conflict" classifiers. "Pure classifiers" were trained to categorize their respective conflict type when the other type was compatible. For example, pure stimulus-conflict classifiers only used ideomotor-compatible stimuli (that were in addition either stimulus-compatible or stimulus-incompatible). "Full classifiers", on the other hand, were trained to categorize their respective conflict type regardless of the compatibility of the other conflict type. For example, a full stimulus-conflict classifier discriminated stimulus-compatible/ideomotor-compatible and stimulus-compatible/ideomotor-incompatible exemplars from stimulus-incompatible/ideomotor-compatible and stimulus-incompatible/ideomotor-incompatible exemplars. These full classifiers resemble the main effects analysis in both the behavioral and univariate fMRI data analyses. The classification accuracies of the 4 types of classifiers were then used in the following categorization scheme that determined domain generality and domain specificity for each searchlight.

Domain Generality/Specificity Classification Analysis
After searchlight MVPA, each searchlight had classification accuracies for pure stimulus-conflict classifiers (r_{ps}), pure ideomotor-conflict classifiers (r_{pi}), full stimulus-conflict classifiers (r_{fs}), and full ideomotor-conflict classifiers (r_{fi}). In order to assess the degree of domain generality/specificity in neural substrates of conflict–control, each searchlight was classified as either domain general, domain specific (stimulus- or ideomotor-conflict-specific), or noninformative.
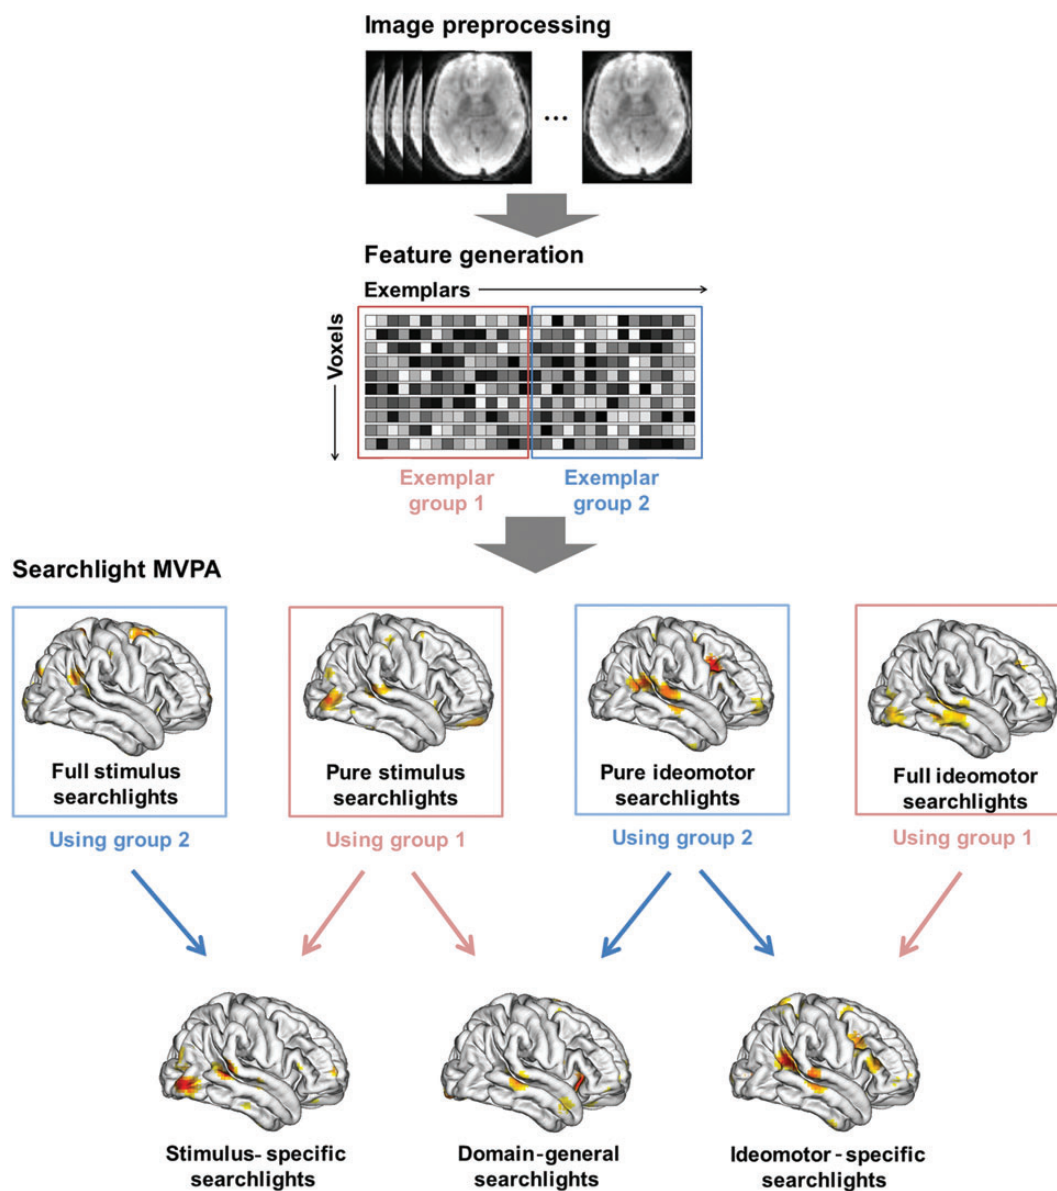
Figure 2. MVPA processing pipeline. Following preprocessing, 24 aggregate exemplars (averaging 4 randomly selected trials) were generated and estimated for each of the 4 experimental conditions (stimulus-compatible/ideomotor-compatible; stimulus-compatible/ideomotor-incompatible; stimulus-incompatible/ideomotor-compatible; stimulus-incompatible/ideomotor-compatible). Then, features from exemplars were split into 2 nonoverlapping groups to avoid potential dependencies in subsequent analyses. For each searchlight, 4 types of independent whole-brain searchlight MVPA classifiers were trained in each subject. "Pure" ideomotor- and stimulus-conflict classifiers were trained to categorize their respective conflict type (i.e. distinguish incompatible from compatible exemplars) when the other conflict type was compatible. "Full" classifiers were trained to categorize their respective conflict type regardless of the compatibility of the other conflict type. Based on these classifiers, domain-general searchlights were defined as searchlights that achieved significant above-chance (corrected for multiple tests at $P < 0.05$) classification accuracy in both types of pure classifiers, whereas domain-specific searchlights were defined as searchlights that could exclusively decode one type of conflict using both pure classifiers and full classifiers.

Specifically, for a given accuracy threshold r, the classification was performed using the following rules: (a) A searchlight is domain general if it can decode both types of conflicts using pure classifiers (i.e. min $(r_{ps}, r_{pi}) \geq r$). (b) A searchlight is domain specific if: It is not domain general ($r_{ps} < r$ or $r_{is} < r$); and it can decode only the stimulus-conflict (i.e. min $(r_{ps}, r_{fs}) \geq r$) or only the ideomotor-conflict (i.e. min $(r_{pi}, r_{fi}) \geq r$) using both pure and full classifiers. (c) A searchlight is noninformative if it is neither domain general nor domain specific.

A number of points regarding these classifications are worth noting: (1) This classification scheme imposes equally strict inclusion thresholds on domain-general and domain-specific searchlights, because both the domain-general and the domain-specific searchlights have to pass 2 independent statistical criteria each. (2) Pure classifiers, rather than full classifiers, were used to test domain generality. This is

because domain-general searchlights may use the same coding for incompatibility regardless of the origins of conflict. Thus, domain-general searchlights may treat both stimulus-compatible/ideomotor-incompatible exemplars, and stimulus-incompatible/ideomotor-compatible exemplars as generic incompatible exemplars. However, in full classifiers, those exemplars would be considered compatible. For example, when training full stimulus-conflict classifiers, stimulus-compatible/ideomotor-incompatible exemplars would be treated as compatible exemplars (because they are both stimulus-compatible). This discrepancy may in theory produce biased results when using full classifiers to test for domain generality. In contrast, using pure classifiers avoids this potential problem by eliminating the above discrepancy. (3) Domain-general and domain-specific searchlights were mutually exclusive because rule (b) required domain-specific searchlights to not be domain-general

searchlights. Furthermore, domain-specific searchlights could be divided into stimulus-conflict-specific and ideomotor-conflict-specific searchlights, based on the conflict they were capable of decoding. These categories were mutually exclusive because stimulus-conflict specificity required $r_{ps} \geq r$ and ideomotor-conflict specificity required $r_{ps} < r$ according to rule (b).

Based on the rules above, we defined $r_g = \min(r_{ps}, r_{pi})$, $r_s = \min(r_{ps}, r_{fs})$, and $r_i = \min(r_{pi}, r_{fi})$ as indices of a searchlight's domain generality, stimulus-conflict-specificity, and ideomotor-conflict-specificity, respectively. In order to avoid bias affecting r_g, r_s, and r_i, it is crucial to keep the pairs of variables in min () independent of each other. However, if r_{ps}, r_{pi}, r_{fs}, and r_{fi} were computed using exemplars from all 4 runs, they could become dependent because they share some exemplars. For example, both r_{ps} and r_{pi} were computed using the same stimulus- and ideomotor- compatible exemplars. This overlap may cause r_{ps} and r_{pi} to be dependent, and thus bias the distribution of r_g. Similarly r_s and r_i could also become biased. To avoid dependence, the exemplars were split into 2 nonoverlapping groups, each of which had exemplars from 2 runs. Then r_{ps} and r_{fi} were estimated using group 1; while r_{pi} and r_{fs} were estimated using group 2. As a result, r_g, r_s, and r_i were computed from independent pairs of variables. With the 4 runs in this task, we used all 6 possible splits (1,2/3,4; 1,3/2,4; 1,4/2,3; 2,3/1,4; 2,4/1,3; and 3,4/1,2). For each split, we computed r_g, r_s, and r_i. We then averaged r_g, r_s, and r_i across splits and used their mean values as the indices of domain generality, stimulus-conflict-specificity, and ideomotor-conflict-specificity, respectively. We refer to r_g, r_s, and r_i as their respective averages across splits.

Nonparametric Statistical Tests of Domain Generality and Domain Specificity

Similar to the situation described above, averaging across splits may cause dependence due to the overlap of exemplars across splits. Furthermore, the dependence is unknown, meaning the null distribution for statistical tests is not available using conventional parametric statistical methods. Thus, we employed a simulation-based, nonparametric method to estimate the null distributions of r_g, r_s, and r_i. Specifically, we simulated features for 21 virtual subjects. Each virtual subject had 96 features (one for each exemplar) on each voxel of the gray matter mask, resembling the structure of features in the original data. Importantly, each feature was sampled from a T-distribution with the same degree of freedom as the original data. This sampling process simulated a situation with no real differences in activation patterns between experimental conditions, making it an ideal candidate for estimating the null distribution. The simulated features then went through searchlight MVPA and the domain generality/specificity classification analysis in exact the same way as the original features and produced simulated r_gs, r_ss, and r_is. The simulated r_gs, r_ss, and r_is at each voxel of each simulated subject were then pooled together to estimate subject-level null distributions of r_g, r_s, and r_i. Using these estimated null distributions, we pursued 2 additional analyses (see next 2 sections), one being a quantification of the incidence of conflict–control domain-general and domain-specific searchlights at the within-subject level, and the other being a group analysis to determine consistently domain-general or -specific brain regions across subjects.

Individual Subject-Level Analysis

The individual subject-level analysis aimed at quantifying the extent of domain generality and domain specificity of cortical conflict–control processing for each subject. This was done by counting the statistically significant domain-general and domain-specific searchlights. To mark potential domain-general and domain-specific searchlights, each subject's r_g, r_s, and r_i maps were transformed to P-value maps based on their corresponding subject-level null distributions. To control for false alarms caused by multiple comparisons, the analysis of functional neuroimages (AFNI) AlphaSim algorithm (Cox 1996) was applied to each subject's P-value maps. Specifically, combined voxel activation intensity and cluster extent thresholds corrected for multiple comparisons were determined by running 5000 Monte Carlo

simulations taking into account the search volume and the estimated smoothness of each axis of the individual SPMs to generate probability estimates of a random field of noise producing a cluster of voxels of a given extent for a set of voxels passing a specific voxel-wise P-value threshold, which we set at 0.05 for all analyses. Given this voxelwise threshold, the simulations determined that cluster sizes >52–62 voxels, depending on the specific analysis, corresponded to combined threshold of $P < 0.05$ (corrected). Note these are comparatively small cluster sizes because the fMRI data were unsmoothed.

In these corrected P-value maps, searchlights within suprathreshold clusters for r_g, r_s, and r_i were counted and used as measures of domain generality, stimulus-conflict-specificity, and ideomotor-conflict-specificity. To further verify the selectivity of domain-specific searchlights, those searchlights were retrained and tested to decode the other type of conflict. In other words, stimulus-conflict-specific searchlights were retrained to decode ideomotor-conflict, and vice versa. In this retrain analysis, full classifiers rather than pure classifiers were used. The reason is that when using domain-specific searchlights to decode their nonspecialized conflict, rule (b) implies that pure classifiers must perform under threshold (see above). Thus, retrained pure classifiers would produce artificially low performance and make them biased for the retrain analysis.

Group-Level Analysis

The group-level analysis aimed at identifying brain regions that were consistently domain general or domain specific across subjects. To test domain generality and domain specificity at the group level using statistical inference, we first estimated the null distributions of group means of r_g, r_s, and r_i by repeatedly sampling from their corresponding subject-level null distributions. Specifically, a number was randomly drawn from each subject-level null distribution. The mean of the 21 numbers was then considered as being sampled from the null distribution of group means. This sampling process was repeated for 10 000 000 times for each subject-level null distribution. The sampled group means were pooled to produce an estimate of the group mean's null distribution. Finally, the newly estimated null distributions were applied to each searchlight to give its group-level P-values of domain generality, stimulus-conflict-specificity, and ideomotor-conflict-specificity. The P-value maps were corrected for multiple comparison using AFNI's AlphaSim algorithm (for details, see above), which determined that voxelwise $P < 0.05$ values paired with cluster sizes of >79–82 voxels, depending on the specific analysis, corresponded to combined multiple comparison-corrected threshold of $P < 0.05$.

Region-of-Interest Analysis

Because the main task was to identify the gender of face stimuli, it was of particular interest to investigate the conflict–control domain specificity of brain regions that specialize in the processing of face stimuli. According to a modular view of conflict–control, only control over stimulus-conflict should involve brain regions representing these stimuli, while control over ideomotor-conflict should not (Egner et al. 2007; Egner 2008). We localized the left (lFFA) and right FFAs (rFFA) using the group-level face > house contrast of a standard GLM analysis. The centers of the region-of-interests (ROIs) were defined as the local maxima of suprathreshold voxels ($P < 0.005$ uncorrected, cluster size >30 voxels) in bilateral fusiform gyrus, and the ROIs themselves were defined as the collection of searchlights whose centers were located within 3 voxels from the center of the ROIs. To perform group-level MVPA analysis within each ROI of each subject, the indices of domain generality, stimulus-conflict-specificity, and ideomotor-conflict-specificity of searchlights were separately z-transformed (i.e. we subtracted the mean and divided by the standard deviation of their corresponding null distributions) and then averaged. These z-scores therefore reflected the relative degree of domain generality/specificity compared with the null distributions. Finally, individual ROI means of z-scores were tested against null (chance-level domain generality/specificity) using group-level, 1-sample t-tests.

Results

Behavioral Data

A repeated-measures ANOVA on mean RT revealed a main effect of stimulus-compatibility ($F_{1,\ 20} = 43.4$, $P < 0.001$), as responses were slower on stimulus-incompatible (696 ms) than stimulus-compatible trials (663 ms), and a main effect of ideomotor compatibility ($F_{1,\ 20} = 27.2$, $P < 0.001$), with slower responses to ideomotor-incompatible (692 ms) than ideomotor-compatible trials (668 ms) (Fig. 1B). Stimulus- and ideomotor-compatibility effects did not interact ($F_{1,\ 20} = 2.5$, n.s.) and their impact on RT was comparable (33 vs. 24 ms, $t_{(20)} = 1.2$, n.s.). Accuracy data conformed to the same pattern of results, displaying noninteracting ($F_{1,\ 20} = 0.3$, n.s.) main effects of stimulus-compatibility ($F_{1,\ 20} = 15.6$, $P < 0.005$) and ideomotor-compatibility ($F_{1,\ 20} = 13.1$, $P < 0.005$), due to more error-prone performance on incompatible than compatible trials (stimulus-compatibility: 95.0% vs. 98.0% accuracy; ideomotor-compatibility: 95.7% vs. 97.4% accuracy).

These results replicate previous findings of independent (additive) conflict–control effects in the stimulus and ideomotor domains (Simon and Berbaum 1990; Kornblum 1994; Egner et al. 2007; Funes et al. 2010) and set the stage for the fMRI data analyses. In principle, the additive behavioral effects could either reflect distinct stages of conflict processing for the 2 domains that are nevertheless resolved by a shared cognitive control resource, or they may reflect the workings of independent, domain-specific conflict–control modules, or a mesh of the two. To adjudicate between these possibilities, we turned to the concurrently acquired fMRI data.
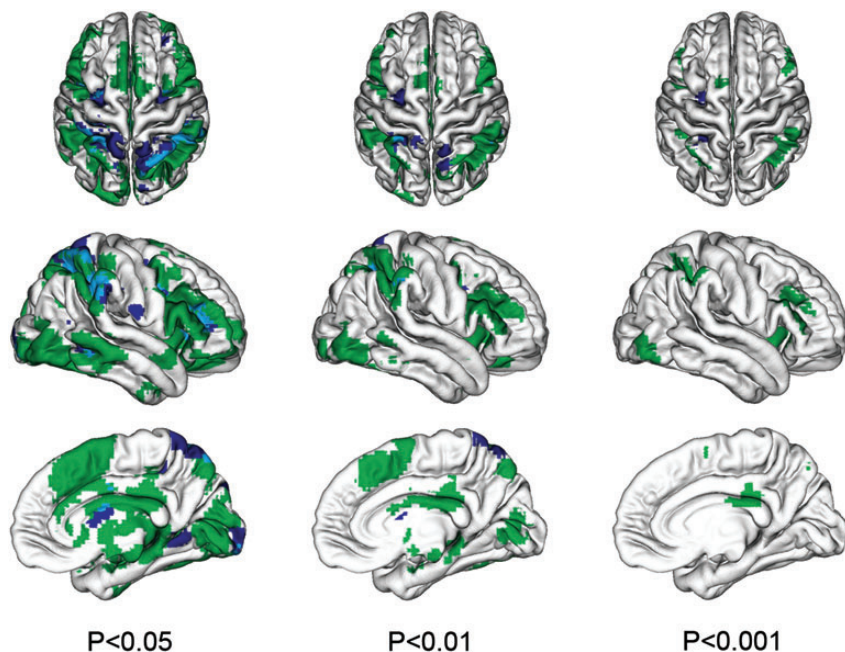
Univariate Imaging Results

For comparison with previous studies and the MVPA analyses (see below), we first conducted conventional (mass-), univariate random-effects group tests on the main effects of stimulus- and ideomotor-compatibility, contrasting activity in incompatible versus compatible trials. The results (Fig. 3 and Table 1) were consistent with the prior literature, as control over stimulus- and ideomotor-conflict elicited activity predominantly in a set of frontal and parietal regions (Peterson et al. 2002; Fan et al. 2003; Liu et al. 2004; Egner et al. 2007), some of which were distinct and some of which overlapped across domains (Peterson et al. 2002; Fan et al. 2003; Liu et al. 2004). Importantly, the degree of this overlap, and thus the implied domain generality of conflict–control, is of course directly dependent on statistical thresholding, in that a more lenient threshold would result in enhanced overlap, while a more stringent threshold would abolish overlap entirely (Fig. 3). Based on the univariate results, one could therefore either conclude that conflict–control processes are implemented in a domain-specific fashion (at more stringent thresholds) or in a domain-general manner (at more lenient thresholds). However, we argue that neither conclusion is warranted, as it is not known whether regions of shared activity carry essential information about the 2 conflict–control processes, nor whether they reflect a functionally homogenous neuronal population response. In addition, key regions of domain-general or domain-specific processing might code information in distributed activity patterns not characterized by mean activation differences between conditions. For a decisive test of domain-general versus domain-specific neural implementation of conflict–control processing, we therefore analyzed the same data set using MVPA.

Within-Subjects MVPA Results

Following whole-brain searchlight MVPA, each searchlight of each subject was associated with indices for domain generality, stimulus-conflict-specificity, and ideomotor-conflict-specificity. These indices were translated into

Figure 3. Univariate fMRI results. Univariate group results for the main effects of stimulus-compatibility (incompatible > compatible, shown in green) and ideomotor compatibility (incompatible > compatible, shown in blue) and their overlap (shown in cyan) are projected onto dorsal, lateral, and medial views of a 3-dimensional-rendered brain (right hemisphere only for lateral and medial views). The maps are displayed at 3 different statistical voxelwise height thresholds ($P < 0.05$, < 0.01, and < 0.001, uncorrected) and a cluster extent threshold of ≥ 10 voxels. For additional details of the activation locations and statistics, see Table 1.

Table 1

Univariate effects of control over stimulus- and ideomotor-conflict

Region	t-value	MNI	Cluster
Main effect of stimulus-conflict (incompatible > compatible)			
Thalamus/caudate	6.76	6, −28, 34	110
Parietal lobule/precuneus	5.80	−30, −58, 55	123
Medial frontal gyrus/cingulate gyrus	5.64	−9, 5, 61	42
Inferior parietal lobule	5.60	−42, −40, 49	69
Middle/inferior frontal gyrus	5.56	45, 41, 25	124
Middle/inferior occipital gyrus	5.06	42, −79, 1	58
Parietal lobule/precuneus	5.00	30, −58, 52	252
Posterior cingulated gyrus	4.60	−24, −76, 7	25
Inferior frontal gyrus	4.53	33, 23, 4	73
Inferior frontal gyrus	4.17	−39, 8, −5	30
Middle/inferior frontal gyrus	4.14	−45, 32, 16	47
Middle/inferior frontal gyrus	4.12	−45, 5, 34	33
Precuneus	4.05	9, −70, 49	10
Middle/inferior temporal gyrus	4.04	−48, −58, −11	19
Main effect of ideomotor-conflict (incompatible > compatible)			
Middle frontal gyrus	5.51	−27, −7, 49	11
Paracentral lobule/precuneus	5.19	−21, −43, 55	14

Note: Inclusion threshold is voxelwise $P < 0.001$ and cluster extent >10.

t-value = t-value at peak voxel, MNI x, y, z = spatial location of peak activation foci in MNI coordinates, cluster = cluster size in number of voxels.

P-values based on simulated null distributions (see Materials and Methods). As a result, each subject had 3 P-value maps, reflecting spatial distributions of P-values for domain generality, stimulus-conflict-specificity, and ideomotor-conflict-specificity. To control for multiple comparisons over searchlights, a whole-brain correction ($P < 0.05$) was applied to the P-maps (see Materials and Methods). Searchlights that passed correction were then considered as either domain general, stimulus-conflict-specific, or ideomotor-conflict-specific at the individual subject level. Many searchlights passed multiple comparison correction. On average of 21 subjects, there were 856.4 ± 225.8 domain-general searchlights, 1344.7 ± 344.0 stimulus-conflict-specific searchlights, and 826.1 ± 198.3 ideomotor-conflict-specific searchlights. A straightforward way of quantifying the degree of domain generality/specificity of the neural organization of conflict–control would be to compare the number of domain-general searchlights with the number of domain-specific searchlights. Yet, direct comparison between the raw numbers of searchlights might be biased by individual differences, because subjects with a larger overall amount of informative searchlights will contribute more strongly to the results. To control for this factor, an analysis was conducted using proportions of searchlights. Specifically, for each subject, we calculated the proportions of domain general, stimulus-conflict-specific, and ideomotor-conflict-specific searchlights relative to all informative searchlights and then compared these proportions across subjects (Fig. 4A). The results showed that the proportion of domain-specific searchlights (mean = 72.7 ± 3.4%) was significantly higher than that of domain-general searchlights (mean = 27.3 ± 3.4%, $t_{(20)} = 5.0$, $P < 0.0001$). Additionally, stimulus-conflict-specific searchlights had a significantly higher proportion (mean = 43.9 ± 5.6%) than domain-general searchlights ($t_{(20)} = 2.2$, $P < 0.05$). The proportion of ideomotor-conflict-specific searchlights (mean = 28.8 ± 5.6%), however, did not differ from domain-general ($t_{(20)} = 1.6$, n.s.) or stimulus-conflict-specific searchlights ($t_{(20)} = 0.2$, n.s.).

Given the preponderance of stimulus-conflict-specific searchlights, we sought to ensure that these data were not a consequence of our whole-brain analysis lacking sensitivity to
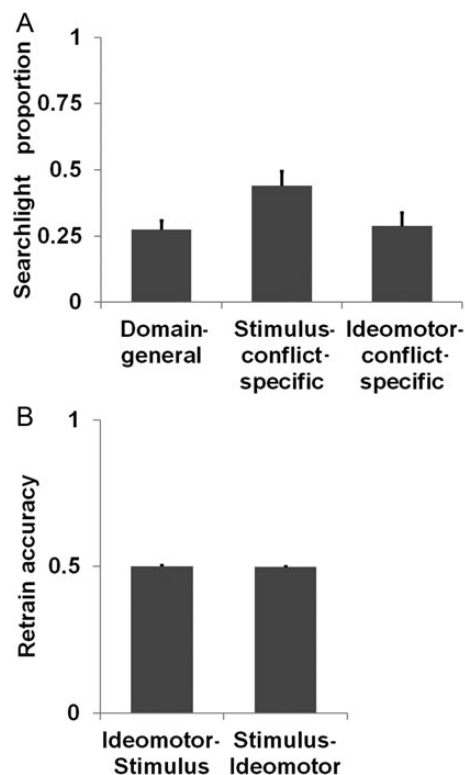
Figure 4. Individual subject-level MVPA results. Domain-general and domain-specific searchlights were identified using nonparametric statistical tests with multiple comparison correction (corrected $P < 0.05$) within each subject. (A) Mean proportion (+MSE) of domain general, stimulus-conflict-specific, and ideomotor-conflict-specific searchlights across subjects. (B) Mean group accuracy (+MSE) of full stimulus-classifiers retrained on ideomotor-conflict-specific searchlights (ideomotor-stimulus) and full ideomotor-classifiers retrained on stimulus-conflict-specific searchlights (stimulus-ideomotor).

both ideomotor- and domain-general conflict–control sites. Therefore, we ran 2 additional analyses where we constrained the MVPA search space to ROIs previously linked to domain-general and ideomotor-conflict-specific control. ROIs were created by combining masks of Brodmann areas (BAs) using Mricron (http://www.mccauslandcenter.sc.edu/mricro/mricron/). In the first analysis, we constrained our within-subject analysis to a putative domain-general cognitive control network described in a recent meta-analysis (Niendam et al. 2012), which consisted of BAs 9, 46, 32, 7, and 40. The results within this ROI showed the same pattern as the whole-brain analysis: Stimulus-conflict-specific searchlights had a significantly higher proportion (mean = 50.6 ± 7.6%) than domain-general searchlights (mean = 23.9 ± 5.7%) ($t_{(20)} = 2.5$, $P < 0.05$). The proportion of ideomotor-conflict-specific searchlights (mean = 25.5 ± 7.4%), however, did not differ from domain-general ($t_{(20)} = 0.2$, n.s.) or stimulus-conflict-specific searchlights ($t_{(20)} = 1.9$, $P < 0.076$). In the second ROI analysis, we constrained our within-subject analysis to premotor and motor regions (i.e. BAs 4 and 6). One subject did not have any searchlight that passed the correction for multiple comparisons and was thus excluded from this analysis. For the remaining 20 subjects, stimulus-conflict-specific searchlights had a significantly higher proportion (mean = 50.0 ± 8.5%) than domain-general searchlights (mean = 12.6 ± 4.4%, $t_{(19)} = 3.6$, $P < 0.005$). The proportion of ideomotor-

conflict-specific searchlights (mean = 37.4 ± 8.6%) was also higher than domain-general searchlights ($t_{(19)}$ = 2.4, $P < 0.05$). There was no significant difference of proportion between stimulus- and ideomotor-conflict-specific searchlights ($t_{(19)}$ = 0.8, n.s.). Thus, unlike at the whole-brain level, ideomotor-conflict-specific searchlights were more prevalent than domain-general searchlights in motor cortices, but not more prevalent than stimulus-conflict-specific searchlights. Taken together, comparisons of proportions of searchlights that were successful at decoding conflict–control processes suggest a preponderance of domain-specific neural substrates of conflict–control at the individual subject level.

To further ascertain the domain specificity of searchlights that were classified as domain-specific, stimulus-conflict-specific searchlights were "retrained" to decode ideomotor-conflict using "full" ideomotor-classifiers, and vice versa. Recall that based on the original classification scheme (see Materials and Methods), stimulus-conflict-specific searchlights had to display suprathreshold accuracy for both pure and full stimulus-conflict-classifiers (rule b), and subthreshold accuracy for pure ideomotor classifiers (rule a). Yet, the accuracy of these searchlights for full ideomotor classifiers remains unknown. Therefore, these searchlights' stimulus-conflict-specificity would be further bolstered if they performed at chance level when retrained with a full ideomotor classifier, and vice versa. As anticipated, domain-specific searchlights showed chance-level retrain performance (Fig. 4B). Specifically, across all 21 subjects, the mean retrain classification accuracy of using stimulus-conflict-specific searchlights to decode ideomotor conflict is 49.8 ± 0.2% ($t_{(20)}$ = −1.2, n.s.). Two subjects had no ideomotor-conflict-specific searchlights after correcting for multiple comparisons. For the remaining 19 subjects, the retrained ideomotor-conflict-specific searchlights also showed chance-level performance when decoding stimulus-conflict (accuracy = 50.1 ± 0.2%, $t_{(18)}$ = −0.6, n.s.). The chance-level retrain performance clearly supports our classification of these searchlights as being purely domain specific rather than being weakly domain general.

Group-Level MVPA Results

The results of the individual subject-level analysis provided evidence for a predominantly domain-specific organization of conflict–control in the human brain, but they granted no information on the anatomical distribution of domain-general and domain-specific conflict–control searchlights. Thus, we further pursued a group-level analysis that aimed at identifying anatomical structures that were consistently domain-general or domain-specific across subjects. Several stimulus-conflict-specific and ideomotor-conflict-specific regions could be identified ($P < 0.05$, whole-brain corrected). Stimulus-conflict-specific searchlights were consistently observed bilaterally in dorsolateral, ventrolateral, and ventromedial prefrontal cortices (PFCs), as well as in parietal and occipito-temporal regions (Fig. 5A and Table 2). Ideomotor-conflict-specific searchlights, on the other hand, were found in the rostral anterior cingulate cortex (ACC) stretching to the head of the caudate, the left lingural gyrus, the left parahippocampal gyrus, and bilateral calcarine sulci (Fig. 5A and Table 2). In contrast, we did not detect any domain-general searchlights that passed whole-brain correction in this group-level analysis, suggesting that domain-general conflict–control processes display more anatomical variability across subjects. We nevertheless pursued an exploratory analysis to provide hints as to where domain-general processing may be observed under higher statistical power (though the results we present necessarily run the risk of entailing false positives). Figure 5B displays group results for domain-general searchlights when the cluster-extent threshold is lowered to >30 voxels (when compared with a threshold of >82 voxels required for whole-brain correction at $P < 0.05$). At this more lenient threshold, potential domain-general conflict control foci emerged in the most
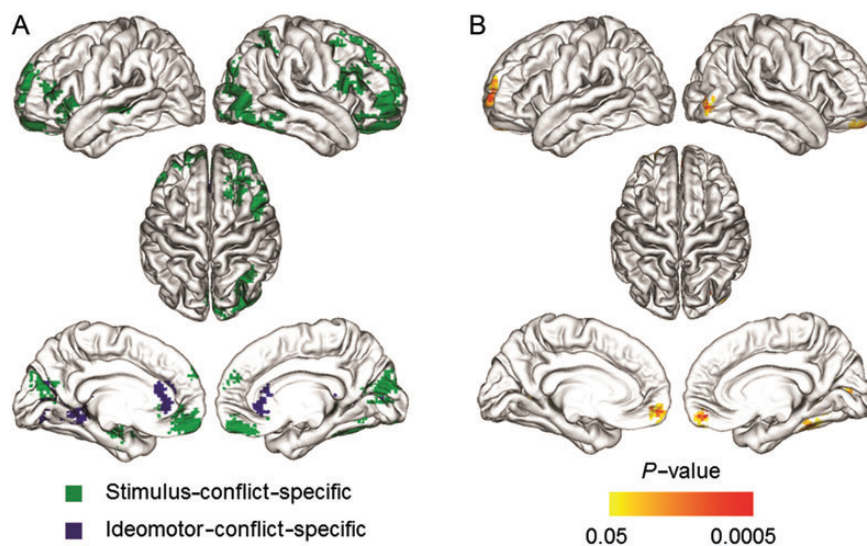
Figure 5. Group-level MVPA results. (A) Searchlights displaying consistent stimulus-conflict-specificity (in green) and ideomotor-conflict-specificity (in blue) across subjects were identified using nonparametric statistical tests and corrected for multiple comparison (corrected $P < 0.05$). No domain-general searchlights were found at this statistical criterion. (B) Exploratory mapping of domain-general searchlights across subjects at a more lenient (not whole-brain corrected) statistical threshold of voxelwise $P < 0.05$ and a cluster extent of 30 voxels.

Table 2
Multivariate effects of control over stimulus- and ideomotor-conflict

Region	t-value	MNI	Cluster
Stimulus-conflict-specific regions			
Orbital gyrus/medial frontal gyrus	5.24	9, 56, −20	343
Hippocampus/parahippocampal gyrus	4.90	−21, −16, −29	97
Superior/middle frontal gyrus	4.63	12, 65, 25	678
Middle occipital gyrus/fusiform gyrus	4.53	30, −82, 28	829
Superior/middle frontal gyrus	4.27	−24, 56, 25	129
Inferior frontal gyrus/insula	3.73	−36, 17, 1	134
Cuneus/precuneus	3.65	6, −82, 31	144
Middle temporal gyrus/putamen	3.38	−36, −19, 1	87
Ideomotor-conflict-specific regions			
Calcarine sulcus	4.94	24, −70, 13	102
Lingural gyrus/parahippocampal gyrus	4.02	−30, −58, −5	131
Calcarine sulcus	3.68	−15, −76, 16	90
Caudate/anterior cingulate cortex	3.14	−21, 23, 4	83

Note: Inclusion threshold is a combined voxel- and cluster-based $P < 0.05$, corrected for multiple comparisons.

t-value = t-value at peak voxel, MNI x, y, z = spatial location of peak activation foci in MNI coordinates, cluster = cluster size in number of voxels.

anterior regions of both lateral and medial PFCs, as well as in extrastriate visual regions of the lateral occipital and the ventral temporal lobes.

Contrasting Univariate and Multivariate Results

It is interesting to note commonalities and differences between the multivariate information-based mapping of conflict-specific control processes (Fig. 5) and the traditional, mass-univariate maps derived from the main effects contrasts of stimulus- and ideomotor compatibility (Fig. 3). The MVPA results concerning regions involved in control over stimulus-conflict are broadly in agreement with the traditional GLM results, although they highlight additional regions involved in stimulus-conflict-specific control that went undetected by the univariate analysis, most notably the ventromedial PFC. Also notable is the absence of the dorsal ACC in the MVPA when compared with the univariate results. The pattern classifier analysis produced more divergent results from the traditional approach with respect to regions involved in control over ideomotor-conflict, where the MVPA results highlight an ideomotor-conflict-specific involvement of the rostral ACC, but discount the importance of dorsal premotor and parietal regions, compared with the mass-univariate results (cf. Figs 5 and 3).

This descriptive comparison across univariate and multivariate analyses provokes the question as to whether the MVPA approach was actually more sensitive in detecting neural substrates of conflict–control than the univariate analysis in the present study. To gauge the unique contribution of MVPA to decoding conflict–control processes, we therefore performed a whole-brain searchlight univariate analysis as comparison. This univariate analysis was the same as the whole-brain MVPA, except that, for each searchlight, the mean of the constituent voxel features was used as the only feature (rather than their multivariate data patterns). After this univariate analysis, the top 1% (~400) of searchlights in terms of the decoding metrics r_{ps}, r_{pi}, r_{fs}, and r_{fi} were averaged to give their corresponding mean of classification for each split and each subject. These mean accuracies reflected the sensitivity of decoding conflict–control in the most informative univariate searchlights. These means were then averaged across
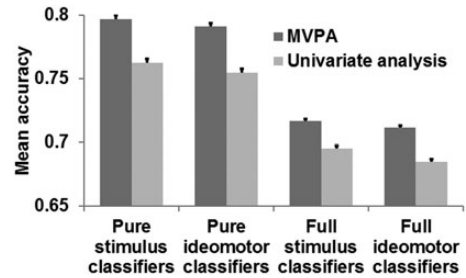

Figure 6. Comparison between univariate and multivariate analyses in decoding conflict–control. Mean classification accuracy (+MSE) of each subject's top 1% accurate searchlights is shown for each type of classifier and each type of analysis (MVPA vs. univariate).

all splits for each subject. For the MVPA results, the individual means were calculated using the same method. Finally, the individual means of r_{ps}, r_{pi}, r_{fs}, and r_{fi} from the univariate analysis results were compared with their counterparts from the MVPA results using group-level, paired t-tests. As shown in Figure 6, compared with the univariate classification results, MVPA results showed significantly higher mean classification accuracy for all 4 types of classifiers (r_{ps}: MVPA mean = 0.797 ± 0.003, univariate analysis mean = 0.763 ± 0.003, $t_{(20)} = 11.3$, $P < 1e-9$; r_{pi}: MVPA mean = 0.791 ± 0.003, univariate analysis mean = 0.755 ± 0.003, $t_{(20)} = 13.5$, $P < 1e-10$; for r_{fs}: MVPA mean = 0.716 ± 0.002, univariate analysis mean = 0.695 ± 0.002, $t_{(20)} = 10.5$, $P < 1e-7$; r_{fi}: MVPA mean = 0.712 ± 0.002, univariate analysis mean = 0.685 ± 0.002, $t_{(20)} = 10.9$, $P < 1e-8$). Additionally, for all types of classifiers and all subjects, mean classification accuracies from MVPA were unanimously higher than their corresponding accuracies from univarate analysis. These results verify unequivocally that MVPA was more sensitive in decoding conflict–control processes than standard, univariate analyses.

ROI-Based Results

The analyses above assessed the degree of domain generality/specificity across the whole brain. In addition, we applied this analysis to brain regions involved in representing the task-relevant stimulus features, that is, face stimuli. This is of particular interest in the present context, because modular models of conflict–control have claimed that control over stimulus-conflict involves biasing the perceptual representation of relevant stimulus features, whereas control over ideomotor-conflict does not (Egner et al. 2007; Egner 2008). Thus, ROIs representing bilateral FFAs were defined using an independent localizer task (Fig. 7A, see Materials and Methods). These ROIs were subsequently tested for domain generality/specificity using ROI mean z-scores at the group level. Data from the left FFA did not allow us to decode conflict–control of any kind above chance level (domain generality: mean = 0.05 ± 0.07, $t_{(20)} = 0.7$, n.s.; stimulus-conflict-specificity: mean = 0.10 ± 0.11, $t_{(20)} = 1.0$, n.s.; ideomotor-conflict-specificity: mean = 0.08 ± 0.07, $t_{(20)} = 1.0$, n.s.). In contrast, the right FFA displayed significant stimulus-conflict-specificity (mean = 0.34 ± 0.09, $t_{(20)} = 3.9$, $P < 0.001$, Fig. 7B), whereas the domain generality (mean = 0.09 ± 0.09, $t_{(20)} = 1.0$, n.s.) and ideomotor-conflict-specificity (mean = −0.05 ± 0.09, $t_{(20)} = 0.5$, n.s.) were at chance level. These results show that, in the present task, where face
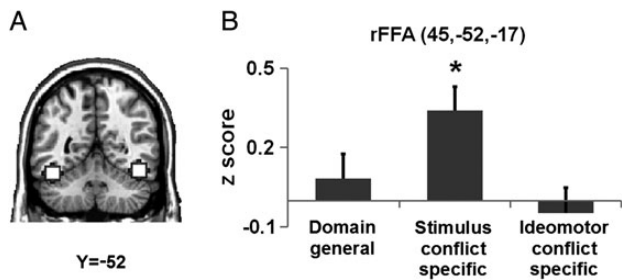
Figure 7. FFA MVPA results. Bilateral group FFA ROIs were defined using an independent localizer task. Within those ROIs, for each subject indices of domain generality, stimulus-conflict-specificity, and ideomotor-conflict-specificity were generated and z-transformed and averaged to produce individual ROI z-scores. (*A*) The spatial location of the group FFA ROIs. (*B*) Mean group z-scores (+MSE) of domain generality, stimulus-conflict-specificity, and ideomotor-conflict-specificity of the right FFA (rFFA, MNI coordinates = [45, −52, −17]). *$P < 0.001$.

stimuli served as task-relevant stimulus features, the (right) FFA is selectively involved in conflict–control in the stimulus domain, but carries no information concerning conflict–control in the ideomotor domain.

Discussion

We harnessed a multivariate, information-based brain mapping approach (Haxby et al. 2001; Haynes and Rees 2006; Kriegeskorte et al. 2006; Norman et al. 2006) to assess the domain specificity of conflict–control. Specifically, we delineated brain regions that carry diagnostic information about stimulus- (or ideomotor-) conflict–control processes, as defined by above-chance success at distinguishing the processing of stimulus- (or ideomotor-) compatible from -incompatible trials and quantified the incidence of regions (searchlights) that were able to decode only stimulus- or ideomotor-conflict–control (domain-specific processors) and those that were able to decode both types of conflict–control processes (domain-general processors). This MVPA approach overcomes several key limitations of the previous, (mass-) univariate fMRI studies in this area (see Introduction).

At the level of behavioral and univariate fMRI data analyses, our study replicated basic findings from the prior literature, producing independent behavioral effects of stimulus- and ideomotor-conflict (Fig. 1*B*), accompanied by (threshold-dependent) distinct as well as shared regional brain activity detected in conventional, univariate fMRI analyses (Fig. 3). Using MVPA, we also established the following, novel findings, which we will discuss in turn: (1) In each subject, many searchlights were observed that could successfully (above chance, and corrected for multiple tests) decode stimulus- and ideomotor-conflict, with a preponderance of domain-specific searchlights, especially those involved in stimulus-conflict–control (Fig. 4). (2) Whereas neural substrates involved in domain-specific conflict–control processes were anatomically consistent across participants, those pertaining to domain-general conflict–control were much less so (Fig. 5). (3) The multivariate information contained in the searchlights was superior in decoding conflict–control than that entailed by univariate, single voxels alone (Fig. 6), and the former identified some brain regions containing information about conflict–control processes that differed from those identified in traditional analyses (cf. Figs 3 and 5). (4) An ROI analysis

focusing on the FFA (as supporting task-relevant stimulus representation) showed that this region was involved in conflict–control over stimulus-conflict, but not in ideomotor-conflict (Fig. 7).

The results of the within-subjects classification analysis suggest that, rather than being entirely domain general or domain specific, conflict–control in the human brain entails both specialized, domain-specific aspects and shared, domain-general processes. This is in line with results from some prior univariate fMRI studies (Milham et al. 2001; Fan et al. 2003; Liu et al. 2004; Kim et al. 2010, 2011), but not with others (Peterson et al. 2002; van Veen and Carter 2005; Liston et al. 2006; Egner et al. 2007). Given the theoretical advantages of the MVPA approach highlighted in the Introduction, and the empirical demonstration in the current paper that the multivariate analysis was superior to the univariate approach in decoding conflict–control processes, we would argue that the present results represent the strongest evidence for a "hybrid" architecture of domain-specific and domain-general conflict–control mechanisms to date.

Within this framework of mixed modular and global conflict–control, however, domain-specific substrates were considerably more pronounced (and anatomically reliable) than domain-general ones. This suggests that distinct conflict–control adaptations have evolved to selectively deal with the challenges of competing representations in different processing pathways. A key advantage that a partly modular or decomposable architecture of control processes confers over centralized control is that it can adapt to changing environmental pressures more flexibly and rapidly than a nonmodular system, while at the same time granting greater robustness against catastrophic systemic dysfunction through mutation, injury, or disease (Simon 1962; Meunier et al. 2010).

Interestingly, in neuroanatomical terms, the mesh of domain-general and domain-specific processors we observed does not necessarily map onto the traditional idea of modular processing in peripheral, sensory cortices combined with central, nonmodular processes in multimodal association cortices (Fodor 1983; Dehaene et al. 1998). Rather, the present data documented a number of domain-specific conflict–control regions even in the PFC, with stimulus-conflict-specific information detected in many regions of lateral PFC, and ideomotor-conflict-specific information in the rostral ACC. Conversely, the exploratory (not whole-brain corrected) group maps of domain-general conflict–control did not exclusively highlight anterior PFC regions, but also posterior regions of the ventral and dorsal visual stream. Taken together, these results argue against a strict mapping of domain-general control processes purely onto fronto-parietal association cortices (Duncan and Owen 2000; Dosenbach et al. 2006; Niendam et al. 2012) and of domain-specific processes solely onto unimodal sensory regions (Fodor 1983).

The specifics of the anatomical distribution of domain-specific and -general searchlights deserve some additional discussion. First, the mere fact that the loci of domain-general processors in our sample appeared to be much less consistent across subjects than those of domain-specific ones may hint at an interesting organizational feature of high-level cognition. Specifically, one could speculate that the very flexibility, which allows certain broad cortical regions to integrate and abstract over multiple sources of information, renders them more heterogeneous in terms of function-to-subregion

mapping than areas concerned only with one particular domain of processing. While our group results on domain-general conflict–control regions have to be interpreted tentatively (due to lenient thresholding), the emergence in particular of frontopolar cortices in this analysis may represent a parallel to previous studies and models that have emphasized a role for regions at the apex of the PFC in the most abstract and integrative forms of controlled processing (Koechlin et al. 2003; Sakai and Passingham 2003; Badre 2008).

Second, it is noteworthy that the MVPA results at times diverged quite markedly from the univariate group maps. The differences observed between the univariate and multivariate group analyses were quite robust to changes in statistical thresholding. Even when running the MVPA group analyses at substantially more lenient cluster-size thresholds, we did not observe substantially greater similarity between the univariate and multivariate group results (data not shown). Notable distinctions with respect to stimulus-conflict-specific foci include the absence of the dorsomedial PFC and dorsal ACC in the MVPA data, and conversely the addition of ventromedial PFC. The former is intriguing vis-à-vis the frequent findings of mean signal enhancements in dorsal ACC in response to conflict (Botvinick et al. 2004; Ridderinkhof et al. 2004) and may cast tentative doubt on whether this region should be thought of as having a specific role in conflict processing (see also Alexander and Brown 2011; Grinband et al. 2011). The latter finding, of ventromedial PFC involvement in conflict–control in the stimulus domain, is also surprising. However, it is plausible that this involvement stems from the fact that the target stimuli in our task were faces. The ventromedial PFC has for some time been thought to belong to a network of regions that contribute to face processing (Summerfield et al. 2006; Ishai 2008), and this area has in fact been found to be highly face sensitive in the FFA localizer in the present paper (data not shown). In line with this interpretation, the present information-based analysis of FFA activity also showed a selective involvement in the control over stimulus-conflict, thus corroborating previous findings and proposals suggesting that this type of conflict–control should indeed involve sensory regions specialized in the processing of the task-relevant stimulus features (Egner and Hirsch 2005; Egner et al. 2007; Egner 2008). Note that it is probable that other brain regions would be involved, in particular in control over stimulus-conflict, if other types of visual (or even auditory) stimuli were employed as target or distracter features.

Thirdly, the group MVPA results concerning key regions bearing information about ideomotor-conflict–control were particularly divergent from the present (and previous) univariate fMRI results: While the latter have typically implicated parietal and premotor regions (e.g. Egner et al. 2007), the present MVPA results highlight primarily the rostral ACC, head of caudate, and ventral visual stream regions, including the parahippocampal gyrus. Given the assumption that ideomotor-conflict should primarily be detected and resolved in regions involved in motor planning and execution (Sturmer et al. 2002; Sturmer and Leuthold 2003; Egner et al. 2007; Egner 2008), these findings are, at first blush, surprising. Closer inspection of the rostral ACC territory implicated in ideomotor-conflict–control in the present study (Fig. 5A), however, suggests that this cluster stretches dorso-caudally into the anterior rostral cingulate motor zone (Picard and Strick 2001), which is thought to be involved in the detection

of response conflict (Botvinick et al. 1999; Picard and Strick 2001). Moreover, the involvement of the rostral ACC in control over ideomotor-conflict has an intriguing precedent in findings from recent lesion studies, where control over ideomotor-conflict (in a Simon task), but not over stimulus-conflict (in a Stroop task), was found to be impaired following damage to the rostral ACC (di Pellegrino et al. 2007; Maier and di Pellegrino, 2012). Similarly, the head of the caudate has been implicated in motor control generally, and in the inhibition of prepotent responses in particular (Vink et al. 2005; Li et al. 2008), an operation that would clearly be highly compatible with the involvement of this region in control over ideomotor-conflict. However, the fact that regions of the ventral and dorsal visual streams also entailed information regarding ideomotor-conflict–control (Fig. 5A) is certainly not in line with a pure motor-system account of how this conflict is generated and resolved.

In closing, one might want to speculate what form the hybrid domain-general/domain-specific conflict–control mechanisms suggested by the present results might take at the level of cognitive or computational architecture. The current data are not incompatible with the basic idea that conflict–control is sub-served by a feedback loop consisting of a mechanism that detects the occurrence of conflict and subsequently triggers the reinforcement of top-down biasing processes comprising the ongoing task set, thus resolving conflict (Botvinick et al. 2001); however, they necessitate the added assumption that multiple, partly dissociable conflict–control loops exist in parallel (Egner 2008). On the other hand, the present data also document some domain-general processing. This may be reflective of some form of central resource, such as attention, that imposes constraints on the kinds and amount of controlled information processing that can be carried out in parallel (Kahneman 1973; Vergauwe et al. 2010). One way to reconcile these findings of partly domain-general and partly (or perhaps mostly) domain-specific resources for conflict–control would be to assume that the detection and resolution of conflict in sensory and motor pathways might be relatively "auxiliary" to this central resource. Corollary data supporting this view comes from studies, showing that adjustments in conflict processing can occur in the absence of awareness both of the conflict (van Gaal et al. 2010) and of the adjustment process itself (Crump et al. 2008).

One possible sketch for such a hybrid executive processing architecture could be that a central, domain-general resource (e.g. attention) is required for setting up and operating a task set (e.g. "categorize faces according to gender; ignore word labels; if male face, press right button; if female face, press left button"), but that the various sensory and motor pathways that are being biased by this task-set are protected by auxiliary conflict adaptation mechanisms that rely largely on domain-specific, peripheral resources. By this view, different conflict–control loops would be anatomically independent and not interfere with each other, as supported in part by the present data set and prior behavioral studies (Simon and Berbaum 1990; Kornblum 1994; Wendt et al. 2006; Egner et al. 2007; Funes et al. 2010). However, they would nevertheless be affected by changes or challenges to the task set itself, such as its reconfiguration during task-switching, or a strong challenge to the task set maintaining central resource pool. There is evidence to support both of these contentions, as

conflict–control has indeed been shown to interact antagonistically with task switching (Goschke 2000; Brown et al. 2007), and to be attenuated by high demands on central processing in dual task settings (Fischer et al. 2008). While concordant with findings in the literature to date, this lose proposal of conflict–control loops as task-set protection mechanisms auxiliary to more central processes that govern task-set content clearly requires future computational elaboration and empirical corroboration.

In conclusion, we employed an information-based brain mapping approach to quantify and map the relative contributions of conflict-specific and conflict-general control mechanisms in the human brain. We document that stimulus- and ideomotor- conflict–control processes can be successfully decoded with neural pattern classifiers, and that the brain regions bearing the discriminating information for this classification are organized primarily in a domain-specific manner, but also include domain-general processors. These data suggest a hybrid architecture of both modular and centralized cognitive control mechanisms.

Funding

Notes

References

Alexander WH, Brown JW. 2011. Medial prefrontal cortex as an action-outomce predictor. Nat Neurosci. 14:1338–1344.

Badre D. 2008. Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. Trends Cogn Sci. 12:193–200.

Botvinick MM, Braver TS, Barch DM, Carter CS, Cohen JD. 2001. Conflict monitoring and cognitive control. Psychol Rev. 108:624–652.

Botvinick MM, Cohen JD, Carter CS. 2004. Conflict monitoring and anterior cingulate cortex: an update. Trends Cogn Sci. 8:539–546.

Botvinick MM, Nystrom LE, Fissell K, Carter CS, Cohen JD. 1999. Conflict monitoring versus selection-for-action in anterior cingulate cortex. Nature. 402:179–181.

Brown JW, Reynolds JR, Braver TS. 2007. A computational model of fractionated conflict-control mechanisms in task-switching. Cognit Psychol. 55:37–85.

Clithero JA, Carter RM, Huettel SA. 2009. Local pattern classification differentiates processes of economic valuation. Neuroimage. 45:1329–1338.

Cox RW. 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Comput Biomed Res. 29:162–173.

Crump MJ, Vaquero JM, Milliken B. 2008. Context-specific learning and control: the roles of awareness, task relevance, and relative salience. Conscious Cogn. 17:22–36.

Dehaene S, Kerszberg M, Changeux JP. 1998. A neuronal model of a global workspace in effortful cognitive tasks. Proc Natl Acad Sci USA. 95:14529–14534.

Desimone R, Duncan J. 1995. Neural mechanisms of selective visual attention. Annu Rev Neurosci. 18:193–222.

di Pellegrino G, Ciaramelli E, Ladavas E. 2007. The regulation of cognitive control following rostral anterior cingulate cortex lesion in humans. J Cogn Neurosci. 19:275–286.

Dosenbach NU, Visscher KM, Palmer ED, Miezin FM, Wenger KK, Kang HC, Burgund ED, Grimes AL, Schlaggar BL, Petersen SE. 2006. A core system for the implementation of task sets. Neuron. 50:799–812.

Duncan J, Owen AM. 2000. Common regions of the human frontal lobe recruited by diverse cognitive demands. Trends Neurosci. 23:475–483.

Egner T. 2007. Congruency sequence effects and cognitive control. Cogn Affect Behav Neurosci. 7:380–390.

Egner T. 2008. Multiple conflict-driven control mechanisms in the human brain. Trends Cogn Sci. 12:374–380.

Egner T, Delano M, Hirsch J. 2007. Separate conflict-specific cognitive control mechanisms in the human brain. Neuroimage. 35:940–948.

Egner T, Etkin A, Gale S, Hirsch J. 2008. Dissociable neural systems resolve conflict from emotional versus nonemotional distracters. Cereb Cortex. 18:1475–1484.

Egner T, Hirsch J. 2005. Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. Nat Neurosci. 8:1784–1790.

Esterman M, Chiu YC, Tamber-Rosenau BJ, Yantis S. 2009. Decoding cognitive control in human parietal cortex. Proc Natl Acad Sci USA. 106:17974–17979.

Fan J, Flombaum JI, McCandliss BD, Thomas KM, Posner MI. 2003. Cognitive and brain consequences of conflict. Neuroimage. 18:42–57.

Fischer R, Dreisbach G, Goschke T. 2008. Context-sensitive adjustments of cognitive control: conflict-adaptation effects are modulated by processing demands of the ongoing task. J Exp Psychol Learn Mem Cogn. 34:712–718.

Fodor JA. 1983. Modularity of mind: an essay on faculty psychology. Cambridge (MA): MIT Press.

Funes MJ, Lupianez J, Humphreys G. 2010. Analyzing the generality of conflict adaptation effects. J Exp Psychol Hum Percept Perform. 36:147–161.

Goschke T. 2000. Intentional reconfiguration and involuntary persistence in task set switching. In: Monsell S, Driver J, editors. Control of cognitive processes: attention and performance XVIII. Cambridge: MIT Press. p. 331–355.

Gratton G, Coles MG, Donchin E. 1992. Optimizing the use of information: strategic control of activation of responses. J Exp Psychol Gen. 121:480–506.

Grinband J, Savitskaya J, Wager TD, Teichert T, Ferrera VP, Hirsch J. 2011. The dorsal medial frontal cortex is sensitive to time on task, not response conflict or error likelihood. Neuroimage. 57:303–311.

Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science. 293:2425–2430.

Haynes JD, Rees G. 2006. Decoding mental states from brain activity in humans. Nat Rev Neurosci. 7:523–534.

Haynes JD, Sakai K, Rees G, Gilbert S, Frith C, Passingham RE. 2007. Reading hidden intentions in the human brain. Curr Biol. 17:323–328.

Howard JD, Plailly J, Grueschow M, Haynes JD, Gottfried JA. 2009. Odor quality coding and categorization in human posterior piriform cortex. Nat Neurosci. 12:932–938.

Ishai A. 2008. Let's face it: it's a cortical network. Neuroimage. 40:415–419.

Kahneman D. 1973. Attention and effort. Englewood Cliffs (NJ): Prentice-Hall.

Kamitani Y, Tong F. 2005. Decoding the visual and subjective contents of the human brain. Nat Neurosci. 8:679–685.

Kanwisher N, McDermott J, Chun MM. 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. J Neurosci. 17:4302–4311.

Kim C, Chongwook C, Jeounghoon K. 2010. Multiple cognitive control mechanisms associated with the nature of conflict. Neurosci Lett. 476:156–160.

Kim C, Kroger JK, Kim J. 2011. A functional dissociation of conflict processing within anterior cingulate cortex. Hum Brain Mapp. 32:304–312.

King JA, Korb FM, von Cramon DY, Ullsperger M. 2010. Post-error behavioral adjustments are facilitated by activation and suppression of task-relevant and task-irrelevant information processing. J Neurosci. 30:12759–12769.

Koechlin E, Ody C, Kouneiher F. 2003. The architecture of cognitive control in the human prefrontal cortex. Science. 302:1181–1185.

Kok P, Jehee JF, De Lange FP. 2012. Less is more: expectation sharpens representastions in the primary visual cortex. Neuron. 75:265–270.

Kornblum S. 1994. The way irrelevant dimensions are processed depends on what they overlap with: the case of Stroop- and Simon-like stimuli. Psychol Res. 56:130–135.

Kornblum S, Hasbroucq T, Osman A. 1990. Dimensional overlap: cognitive basis for stimulus-response compatibility—a model and taxonomy. Psychol Rev. 97:253–270.

Kriegeskorte N, Goebel R, Bandettini P. 2006. Information-based functional brain mapping. Proc Natl Acad Sci USA. 103:3863–3868.

Li CS, Yan P, Sinha R, Lee TW. 2008. Subcortical processes of motor response inhibition during a stop signal task. Neuroimage. 41:1352–1363.

Liston C, Matalon S, Hare TA, Davidson MC, Casey BJ. 2006. Anterior cingulate and posterior parietal cortices are sensitive to dissociable forms of conflict in a task-switching paradigm. Neuron. 50:643–653.

Liu X, Banich MT, Jacobson BL, Tanabe JL. 2004. Common and distinct neural substrates of attentional control in an integrated Simon and spatial Stroop task as assessed by event-related fMRI. Neuroimage. 22:1097–1106.

Maier ME, di Pellegrino G. 2012. Impaired conflict adaptation in an emotional task context following rostral anterior cingulate cortex lesions in humans. J Cogn Neurosci. 24:2070–2079.

Meunier D, Lambiotte R, Bullmore ET. 2010. Modular and hierarchically modular organization of brain networks. Front Neurosci. 4:200.

Milham MP, Banich MT, Webb A, Barad V, Cohen NJ, Wszalek T, Kramer AF. 2001. The relative involvement of anterior cingulate and prefrontal cortex in attentional control depends on nature of conflict. Brain Res Cogn Brain Res. 12:467–473.

Miller EK, Cohen JD. 2001. An integrative theory of prefrontal cortex function. Annu Rev Neurosci. 24:167–202.

Misaki M, Kim Y, Bandettini PA, Kriegeskorte N. 2010. Comparison of multivariate classifiers and response normalizations for pattern-information fMRI. Neuroimage. 53:103–118.

Niendam TA, Laird AR, Ray KL, Dean YM, Glahn DC, Carter CS. 2012. Meta-analytic evidence for a superordinate cognitive control network subserving diverse executive functions. Cogn Affect Behav Neurosci. 12:241–268.

Norman DA, Shallice T. 1986. Attention to action: willed and automatic control of behavior. In: Schwarz GE, Shapiro D, editors. Consciousness and self-regulation. New York: Plenum Press.

Norman KA, Polyn SM, Detre GJ, Haxby JV. 2006. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. Trends Cogn Sci. 10:424–430.

Peterson BS, Kane MJ, Alexander GM, Lacadie C, Skudlarski P, Leung HC, May J, Gore JC. 2002. An event-related functional MRI study comparing interference effects in the Simon and Stroop tasks. Brain Res Cogn Brain Res. 13:427–440.

Picard N, Strick PL. 2001. Imaging the premotor areas. Curr Opin Neurobiol. 11:663–672.

Ridderinkhof KR, Ullsperger M, Crone EA, Nieuwenhuis S. 2004. The role of the medial frontal cortex in cognitive control. Science. 306:443–447.

Sakai K, Passingham RE. 2003. Prefrontal interactions reflect future task operations. Nat Neurosci. 6:75–81.

Shin YK, Proctor RW, Capaldi EJ. 2010. A review of contemporary ideomotor theory. Psychol Bull. 136:943–974.

Simon HA. 1962. The architecture of complexity. Proc Am Philos Soc. 106:467–482.

Simon JR. 1969. Reactions toward the source of stimulation. J Exp Psychol. 81:174–176.

Simon JR, Berbaum K. 1990. Effect of conflicting cues on information processing: the "Stroop effect" vs. the "Simon effect". Acta Psychol (Amst). 73:159–170.

Stroop JR. 1935. Studies of interference in serial verbal reactions. J Exp Psychol. 18:643–662.

Sturmer B, Leuthold H. 2003. Control over response priming in visuomotor processing: a lateralized event-related potential study. Exp Brain Res. 153:35–44.

Sturmer B, Leuthold H, Soetens E, Schroter H, Sommer W. 2002. Control over location-based response activation in the Simon task: behavioral and electrophysiological evidence. J Exp Psychol Hum Percept Perform. 28:1345–1363.

Summerfield C, Egner T, Greene M, Koechlin E, Mangels J, Hirsch J. 2006. Predictive codes for forthcoming perception in the frontal cortex. Science. 314:1311–1314.

van Gaal S, Lamme VA, Ridderinkhof KR. 2010. Unconsciously triggered conflict adaptation. PLoS One. 5:e11508.

van Veen V, Carter CS. 2005. Separating semantic conflict and response conflict in the Stroop task: a functional MRI study. Neuroimage. 27:497–504.

Vergauwe E, Barrouillet P, Camos V. 2010. Do mental processes share a domain-general resource? Psychol Sci. 21:384–390.

Vink M, Kahn RS, Raemaekers M, van den Heuvel M, Boersma M, Ramsey NF. 2005. Function of striatum beyond inhibition and execution of motor responses. Hum Brain Mapp. 25:336–344.

Wendt M, Kluwe RH, Peters A. 2006. Sequential modulations of interference evoked by processing task-irrelevant stimulus features. J Exp Psychol Hum Percept Perform. 32:644–667.