

Tracking cognitive fluctuations with multivoxel pattern time course (MVPTC) analysis

Yu-Chin Chiu^{*,1}, Michael S. Esterman¹, Leon Gmeindl, Steven Yantis

Department of Psychological and Brain Sciences, Johns Hopkins University, United States

ARTICLE INFO

Article history:

Received 8 February 2011

Received in revised form 9 July 2011

Accepted 9 July 2011

Available online 23 July 2011

Keywords:

Attention
Cognitive control
Basal ganglia
Switching
MVPA
fMRI
SVM

ABSTRACT

The posterior parietal cortex, including the medial superior parietal lobule (mSPL), becomes transiently more active during acts of cognitive control in a wide range of domains, including shifts of spatial and nonspatial visual attention, shifts between working memory representations, and shifts between categorization rules. Furthermore, spatial patterns of activity *within* mSPL, identified using multivoxel pattern analysis (MVPA), reliably distinguish between different acts of control. Here we describe a novel multivoxel pattern-based analysis that uses fluctuations in cognitive state over time to reveal inter-regional functional connectivity. First, we used MVPA to model patterns of activity in mSPL associated with shifting or maintaining spatial attention. We then computed a multivoxel pattern time course (MVPTC) that reflects, moment-by-moment, the degree to which the pattern of activity in mSPL more closely matches an attention-shift pattern or a sustained-attention pattern. We then entered the MVPTC as a regressor in a univariate (i.e., voxelwise) general linear model (GLM) to identify voxels whose BOLD activity covaried with the MVPTC. This analysis revealed several regions, including the striatum of the basal ganglia and bilateral middle frontal gyrus, whose activity was significantly correlated with the MVPTC in mSPL. For comparison, we also conducted a conventional functional connectivity analysis, entering the mean BOLD time course in mSPL as a regressor in a univariate GLM. The latter analysis revealed correlations in extensive regions of the frontal lobes but not in any subcortical area. The MVPTC analysis provides greater sensitivity (e.g., revealing the striatal-mSPL connectivity) and greater specificity (i.e., revealing more-focal clusters) than a conventional functional connectivity analysis. We discuss the broad applicability of MVPTC analysis to a variety of neuroimaging contexts.

Published by Elsevier Ltd.

1. Introduction

Humans prioritize the processing of goal-relevant sensory information through voluntary shifts of selective attention. Several recent studies have reported that voluntary shifts of attention are associated with transient activity in the medial wall of the superior parietal lobule (mSPL; e.g., Kelley, Serences, Giesbrecht, & Yantis, 2008; Liu, Slotnick, Serences, & Yantis, 2003; Serences, Schwarzbach, Courtney, Golay, & Yantis, 2004; Shomstein & Yantis, 2006; Shulman et al., 2009; Vandenberghe, Gitelman, Parrish, & Mesulam, 2001; Yantis et al., 2002). In contrast, sustained activity associated with maintaining attention is found in intraparietal sulcus and prefrontal cortex (e.g., Saygin & Sereno, 2008; Serences & Yantis, 2007; Silver, Ress, & Heeger, 2005). These findings sug-

gest that mSPL plays a role in reconfiguring or shifting attention rather than in maintaining the current state of attention. Furthermore, this transient mSPL signal is also observed in non-spatial and non-perceptual acts of control, including shifts of categorization rule (Chiu & Yantis, 2009) and shifts of attention between working memory representations (Tamber-Rosenau, Esterman, Chiu, & Yantis, 2011). These studies thus implicate mSPL as a domain-independent hub for cognitive reconfiguration (Chiu & Yantis, 2009; Greenberg, Esterman, Wilson, Serences, & Yantis, 2010).

In addition to mSPL, other cortical and subcortical regions have been associated with cognitive reconfiguration. For example, dorsolateral prefrontal regions are often co-activated with parietal cortex, forming a dorsal frontoparietal attention control network for deployments of attention to goal-relevant sensory information (Corbetta & Shulman, 2002). Furthermore, many studies have demonstrated functional connectivity between dorsal parietal cortex and prefrontal cortex (e.g., Fox, Corbetta, Snyder, Vincent, & Raichle, 2006), and have shown that this connectivity affected behavior during spatial attention tasks (He, Shulman, Snyder, & Corbetta, 2007; Thiebaut de Schotten et al., 2005). In addition to those cortical regions, the basal ganglia (BG) also have been

* Corresponding author at: Department of Psychology, University of California, San Diego, 9500 Gilman Drive, #0109, La Jolla, CA 92093-0109, United States. Tel.: +1 443 3702325.

E-mail address: chiu.yuchin@gmail.com (Y.-C. Chiu).

¹ These authors contributed equally to this work.

implicated in shifts of spatial attention (e.g., Gitelman et al., 1999; Grande et al., 2006; Shulman et al., 2009), shifts of task set (e.g., Cools, Altamirano, & D'Esposito, 2006; Cools, Clark, & Robbins, 2004; Leber, Turk-Browne, & Chun, 2008; Ravizza & Ivry, 2001), and updates in working memory (O'Reilly & Frank, 2006). However, functional connectivity has not been demonstrated between cortical control regions and subcortical structures (e.g., the BG) in humans during shifts of spatial attention. The current study was designed to examine the functional connectivity between mSPL and other cortical and subcortical regions using a novel multivariate technique.

One well-established approach to assessing functional connectivity first computes the time course of the mean blood oxygenation level dependent (BOLD) signal across all voxels within a seed region of interest (ROI). This time course is then entered as a univariate (voxelwise) general linear model (GLM) regressor in order to identify other voxels whose activity covaries with that of the seed region (e.g., Biswal, Yetkin, Haughton, & Hyde, 1995; Friston, Frith, Liddle, & Frackowiak, 1993). Another approach, beta series correlation (Rissman, Gazzaley, & D'Esposito, 2004), uses trial-by-trial beta coefficient values, instead of raw time series, within a seed region to explore correlations across the brain. More sophisticated approaches, such as psychophysiological interaction analysis (e.g., Diekhof, Falkai, & Gruber, 2009; Duann, Ide, Luo, & Li, 2009) or dynamic causal modeling (e.g., Friston et al., 1997; Smith, Stephan, Rugg, & Dolan, 2006; Stephan, Marshall, Penny, Friston, & Fink, 2007), further include interaction terms among regressors to explore connectivity. Using the mean BOLD signal (or beta values) within an ROI to explore functional connectivity, however, assumes that all voxels within the ROI are all estimates of a single, common time series (thus justifying taking their average). Furthermore, this method relies on the presence of reliable activations (after correcting for multiple comparisons) in the seed regions of interest.

Here we develop an *information-based* functional connectivity method. In general, an information-based approach employs multivoxel pattern analysis (MVPA) to identify regions in which information expressed by spatiotemporal patterns of voxel activation reliably reflects distinct cognitive states (e.g., Chiu, Esterman, Han, Rosen, & Yantis, 2011; Kamitani & Tong, 2005; Kriegeskorte, Goebel, & Bandettini, 2006; Norman, Polyn, Detre, & Haxby, 2006; Polyn, Natu, Cohen, & Norman, 2005; Serences, Ester, Vogel, & Awh, 2009). Our novel information-based functional connectivity method exploits the multivoxel pattern of activity within a seed region by computing a continuous index of "pattern strength" as it evolves and fluctuates over time. This multivoxel pattern time course (MVPTC) then is used to identify other voxels whose activity covaries with the pattern strength within the seed region. This approach was motivated by recent studies (Esterman, Chiu, Tamber-Rosenau, & Yantis, 2009; Greenberg et al., 2010; Tamber-Rosenau et al., 2011) showing that spatial patterns of activity within mSPL reliably predicted which of several domains of cognitive control subjects were engaged in (e.g., spatial representations vs. rule representations; and visual vs. working memory representations) at each time point of an experimental run. Thus, multivoxel patterns of activity in mSPL provide a signature of cognitive state that is otherwise masked by averaging the BOLD signal across voxels within the ROI.

In the current study, subjects performed cue-evoked covert shifts of spatial attention in the task. We first identified a region of mSPL that exhibited transient increases in mean activity time-locked to shifts of attention. We then used MVPA (employing a linear support vector machine, or SVM) to model multivoxel activation patterns within mSPL associated with attention shifts and sustained attention, respectively. The resulting classifier was then applied to each time point in the entire mSPL time series, yielding

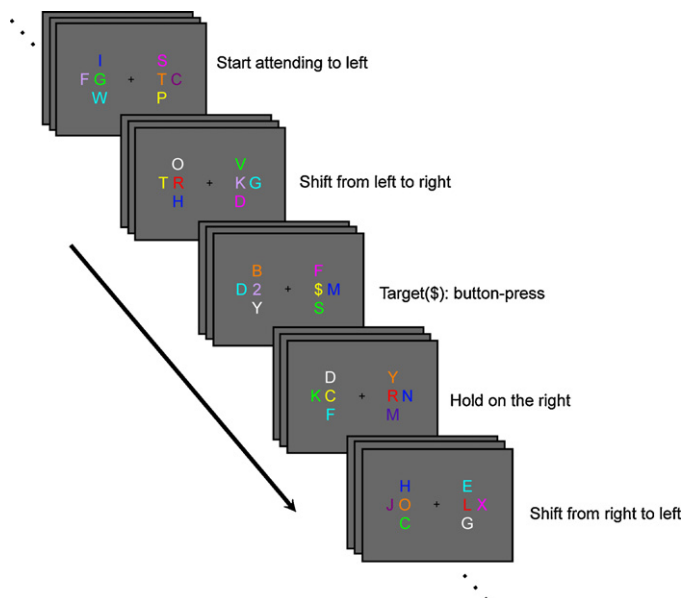


Fig. 1. Behavioral task. While maintaining fixation at the center of the screen, the subject was instructed to start attending to the left-central RSVP stream and to monitor for occasional targets (2, 3, #, \$). When a shift cue was presented in the left-central stream, the subject had to shift attention covertly to the right-central RSVP stream, and vice versa, and continue to monitor for targets in the currently relevant stream. The subject had to press one of four buttons to indicate the identity of each target appearing in the relevant stream.

a continuous information-based index (i.e., the classifier decision value) of the degree to which the mSPL pattern reflected shifting vs. sustained attention on a moment-by-moment basis. Finally, the MVPTC obtained from mSPL was entered as a regressor in a whole-brain voxelwise GLM to explore information-based functional connectivity between mSPL and the rest of the brain.

2. Methods


2.1. Subjects

Ten right-handed subjects (20–23 yrs old, four males) participated in this study. Two of the subjects failed to maintain fixation during the task and were excluded from further analysis. All subjects had normal or corrected-to-normal vision and had no history of neurological impairment. The Johns Hopkins Medicine Institutional Review Board approved the study protocol. Written informed consents were obtained from all participants.


2.2. Stimuli and procedure

Subjects were instructed to maintain fixation on a central fixation cross while monitoring one of eight rapid serial visual presentation (RSVP) streams of alphanumeric characters (see Fig. 1). Each RSVP frame had a duration of 125 ms. Two of these streams were task-relevant, located 3.5° to the left and right of the fixation cross on the horizontal meridian; each was flanked by three irrelevant distractor streams (3.3° center-to-center distance) in order to maximize the demand for selective attention. Critical events (cues and targets) were randomly intermixed among filler items (A–Z, except for R and L) within the task-relevant streams; the distractor streams contained only filler items. All characters subtended approximately 1.4° in height and 1.0° in width, and were presented in uppercase Arial font. Cues were always rendered in bright red (and therefore were highly salient), while all other items (targets and fillers) were randomly rendered in one of eight colors (green, blue, purple, orange, yellow, army green, rose, and cyan). No two items had the same color in any given frame. Subjects began with their attention directed to one of the two relevant streams (according to instructions at the beginning of each run). 'Shift' and 'Hold' instructions were conveyed by the letters "L" and "R" in the currently attended stream. These cues instructed subjects to direct attention to the left or right stream, respectively (e.g., "R" in the left stream signaled subjects to shift attention to the right stream; "L" in the left stream signaled subjects to maintain attention to the left stream). Cue onset asynchrony varied between 10 and 15 s, in 1-s increments, and the minimum target-to-cue and cue-to-target onset asynchrony was 3 s. Approximately half of the cues were shift cues, and half were hold cues.

A. Define ROI: shift vs. hold contrast



B. Extract multivoxel patterns from ROI C. Classify multivoxel patterns using SVM



D. Index multivoxel pattern time course

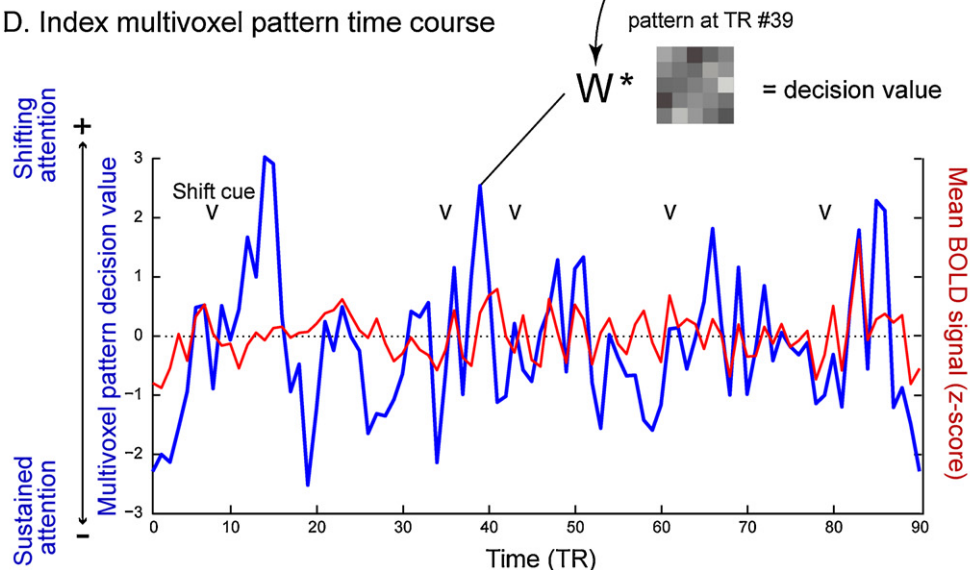


Fig. 2. Multivoxel pattern time course (MVPTC) analysis flow. (A) Shift-related voxels in mSPL as identified by the HRF-based univariate GLM. (B) A schematic illustration of an event-related average BOLD time course in mSPL time-locked to shift cues. The dashed rectangles indicate the time windows that were used to extract multivoxel patterns associated with sustained attention and shifts of attention, respectively, during classifier training. (C) An SVM classifier was trained to discriminate patterns associated with shifts of attention vs. sustained attention and tested using a leave-one-run-out protocol. (D) An example of an mSPL MVPTC (the SVM multivoxel pattern decision value at each time point, in blue) and mSPL mean BOLD time course (expressed as a z-score, in red) from a single run of one representative participant. Arrows mark the onset of shift cues (e.g., sLR or sRL).

Subjects were instructed to make a four-alternative button-press response to each appearance of a target character (2, 3, #, or \$) embedded within the currently relevant RSVP stream. Two target characters were simultaneously presented; either a “2” or “3” was presented in one stream and either a “#” or “\$” was presented in the other stream. Two stimulus-response mappings were counterbalanced across subjects: 2, 3, #, \$, or #, \$, 2, 3 were mapped to the right index-, middle-, ring-, and little-finger buttons, respectively. Behavioral responses thus provided a check on the subject’s current locus of attention. Note that the behavioral responses always occurred during epochs of sustained attention; no behavioral response was made to the cues themselves. Each subject completed 10–13 runs, each lasting 134 s within one 2-h fMRI session.

Prior to the fMRI session, subjects completed a behavioral practice session. During this session, the rate of stimulus presentation gradually decreased from 500 to

125 ms per frame, a rate at which subjects were able to maintain a performance accuracy of at least 80% across two successive blocks. In the practice session, accuracy feedback was provided at the end of each block. In the imaging session, frame duration was fixed at 125 ms and accuracy feedback was not provided.

During scanning, visual stimuli were projected onto a screen placed at the end of the magnet bore and viewed with a mirror mounted above the head coil from a distance of 68 cm. Subjects held a custom-built MR-compatible response box with four buttons in their right hand. All visual stimuli were presented on a grey background. Stimulus presentation and behavioral data collection were controlled by custom MATLAB (The MathWorks, Natick, MA) code using Psychophysics Toolbox 3 (Brainard, 1997; Pelli, 1997). Eye position was monitored in the imaging session using a custom MR-compatible infrared camera and ViewPoint 2.8.3 eye tracking software (Arrington Research, Scottsdale, AZ).

2.3. Image acquisition and analysis

Neuroimaging data were acquired with a Philips Intera 3T scanner and an 8-channel SENSE head coil (MRI Devices) at the F.M. Kirby Research Center for Functional Brain Imaging (Baltimore, MD). One high-resolution, whole-brain anatomical scan was acquired with an MPRAGE T1-weighted sequence yielding 1-mm isotropic voxel resolution (coronal slices, matrix = 256 × 256, TE = 3.7 ms, TR = 8.0 ms, flip angle = 8°). Whole-brain functional volumes were acquired with a T2*-weighted echoplanar imaging (EPI) sequence in 31 axial slices (3-mm thickness, 1-mm gap, 3 mm × 3 mm in-plane resolution, matrix = 64 × 64, TE = 30 ms, TR = 1.5 s, flip angle = 70°).

Neuroimaging data were analyzed using BrainVoyager QX software (Brain Innovation, The Netherlands). Functional data were slice-time and motion corrected and then temporally high-pass filtered to remove components occurring three or fewer cycles per run. To correct for between-scan motion, each subject's EPI volumes were all coregistered to that subject's anatomical scan. Finally, the volumes were Talairach-transformed and resampled into 3-mm isotropic voxels. No spatial smoothing was applied to the EPI data.

2.4. Functional region of interest (functional ROI)

The GLM approach (Friston et al., 1995) was used to estimate beta weights for regressors representing critical events. Included were regressors that modeled the presentation of targets and cues instructing subjects to shift attention from left to right (sLR), shift attention from right to left (sRL), hold attention on the right stream (hR), and hold attention on the left stream (hL), respectively. The regressors were created by convolving a single-gamma hemodynamic response function (Boynton, Engel, Glover, & Heeger, 1996) with a stick function marking the onset of each event. The mSPL ROI to be used for pattern classification was identified by a contrast of attention Shift (sLR, sRL) vs. attention Hold (hR, hL) with a liberal threshold of $t(7) = 2$ ($p < 0.06$ uncorrected; see Fig. 2A). The mSPL ROI in Talairach space was then transformed back to the subjects' native anatomical space, from which preprocessed functional EPI data were extracted for further analyses. The BOLD signal was z-transformed relative to the mean and standard deviation of each run.

2.5. Multivoxel pattern classification

Patterns associated with shifting vs. sustained attention were extracted from the mSPL ROI for each subject: for each voxel in the mSPL ROI, activity associated with shifting attention was defined as the mean of the BOLD signal at 6 and 7.5 s after the onset of an attention-shift cue. Activity associated with sustained attention was defined as the mean of the BOLD signal at 1.5 and 0 s before the onset of an attention-shift cue, an epoch during which participants should have sustained attention to the left or right RSVP stream (see Fig. 2B). It is possible that this pre-shift epoch included other cognitive operations or states in addition to sustained attention (e.g., letter identification). However, it should be noted that discriminating the target stimuli in the RSVP task required highly focused attention, and subjects did so with reasonable accuracy.

This resulted in a set of 120–128 training/testing examples for each pattern for each subject. A linear classifier was trained to discriminate between these multivoxel patterns associated with shifting vs. sustained attention for each subject. We used the standard leave-one-run-out cross-validation procedure (see Fig. 2C) with a support vector machine (OSU SVM toolbox, an adaptation of libsvm, Chang & Lin, 2001). The initial mean cross-validation accuracy at the group level was 65.1% (S.E.M. = 1.95%), which was significantly above chance ($p < 0.05$), indicating that patterns of activity in mSPL can be reliably decoded when all of the voxels within the ROI were used. This was, of course, as expected, because several previous studies have found transient activity in mSPL during shifts of attention, but not during sustained state of attention (i.e., Kelley et al., 2008; Saygin & Sereno, 2008; Silver et al., 2005; Yantis et al., 2002). Nevertheless, to ensure that an optimized classifier would be employed when distinguishing patterns of activity at other time points not included in the training and testing epochs above, we eliminated voxels within the ROI that contributed mostly noise (i.e., voxels that did not help to distinguish between shifting and sustained attention patterns), as described below.

For each fold of training and testing, a set of classifier weights (one weight per voxel) was obtained. We first calculated the absolute value of the mean weight for each voxel across all folds, and then ranked them. Based on this ranking, classifier accuracy was optimized with the 75 mSPL voxels with the highest importance (i.e., largest absolute weight magnitude; see Esterman et al., 2009). The selected voxels were then re-trained with all examples (no run left-out) to obtain a new set of weights. Note that no inference can be made from this analysis about the voxels in mSPL (Kriegeskorte, Simmons, Bellgowan, & Baker, 2009).

2.6. Multivoxel pattern-based functional connectivity

The classifier was then applied to the mSPL activity pattern at each time point of all runs to derive a decision value (see Fig. 2D), thus creating a multivoxel pattern time course, or MVPTC. The MVPTC indexes the degree to which activation in the ROI reflects a shifting-attention pattern (i.e., a positive decision value) vs. a sustained-attention pattern (i.e., a negative decision value) as multivoxel activity

fluctuates throughout task performance. Importantly, the output of the classifier was not binarized (as it otherwise would be for discrete classification; see Fig. 2D for an example).

The MVPTC for each run was entered as a regressor in a whole-brain univariate (voxelwise) GLM; we refer to this as the *mSPL-pattern* model. To account for whole-brain fluctuations in signal intensity (Fox, Zhang, Snyder, & Raichle, 2009; Zarahn, Aguirre, & D'Esposito, 1997), we also entered a global signal regressor (mean activity across all voxels) in this GLM as a nuisance regressor. A random-effect group level analysis (with subject as a random factor and run number as a nuisance variable) was conducted to identify voxels that exhibited activity that was significantly correlated with the mSPL MVPTC. Because the mSPL MVPTC was not independent of voxel activity in mSPL itself and the surrounding parietal regions (see Kriegeskorte et al., 2009; Vul, Harris, Winkelman, & Pashler, 2009), voxels in the parietal lobe were excluded from the results reported below.

For comparison with conventional functional connectivity analysis, we generated a regressor in which the BOLD signal across the same 75 mSPL voxels was averaged to create a mean time course (termed the *mSPL-mean* model, see Fig. 2D for an example). This was equivalent to assuming that all voxels contribute equally by setting identical classifier weight for each voxel. The mSPL-mean model then served as a regressor in a separate whole-brain GLM along with the global signal regressor as described above. A random-effect group level analysis (with subject as a random factor and run number as a nuisance variable) was conducted to identify voxels that were significantly correlated with the mean activity in mSPL. Statistical maps were corrected for multiple comparisons by applying a cluster-size threshold (Forman et al., 1995): voxelwise nominal $p = 0.009$, $t(7) = 3.5$, corrected $\alpha = 0.05$.

3. Results

3.1. Behavioral results

During the RSVP task in the scanner, mean accuracy was 83.3% (SE = 2.5%). Almost all of the errors were misses (i.e., failure to respond; $M = 15.2\%$, SE = 2.4%) rather than incorrect button-presses ($M = 1.5\%$, SE = 0.5%), indicating that subjects rarely failed to shift attention to the appropriate RSVP stream following a shift cue (i.e., if subjects failed to shift attention to the relevant stream, they would have responded often to targets in the irrelevant stream, resulting in incorrect button-presses).

3.2. Neuroimaging results

The Shift vs. Hold contrast in the hemodynamic response function (HRF)-based univariate GLM yielded a region in mSPL that was subsequently used as the seed region for connectivity analyses (Fig. 2A). No other brain region exceeded the threshold for statistical significance. Fig. 2D shows the MVPTC (blue) and the mean BOLD signal (red) from mSPL for a representative subject during one run. The small arrows indicate when a shift-cue appeared—that is, when the subject was instructed to shift attention. Both time courses exhibit fluctuations; in most instances, the MVPTC increases in magnitude sharply after shift cues, but to a varying degree from one shift cue to the next. The mean BOLD signal also increases following shifts but the change in magnitude is less pronounced (see Fig. 2D). Using this seed region, we computed a whole-brain correlation map with the standard functional connectivity analysis approach (i.e., the mSPL-mean model; Fig. 3A). As shown in Fig. 3A, the mSPL-mean model identified a network of frontal and occipital regions; see Table 1 for a complete list.

After computing the MVPTC from this same region, several clusters of voxels were found to exhibit reliable positive correlations (i.e., correlations with shift-like multivoxel activity in mSPL) including the left caudate nucleus of the basal ganglia (BG) and bilateral middle frontal gyrus (MFG; see Fig. 3B). We found no brain region that exhibited a negative correlation with the MVPTC. See Table 1 for a complete list.

The HRF-based univariate GLM (Fig. 2A) failed to reveal frontal or BG activations associated with shifts of attention, when examining either a Shift vs. Hold contrast or a simple main effect of Shift (even with a liberal threshold of $t(7) = 2.3$, voxelwise $p < 0.06$, uncorrected). Although the MFG was identified in the mSPL-mean

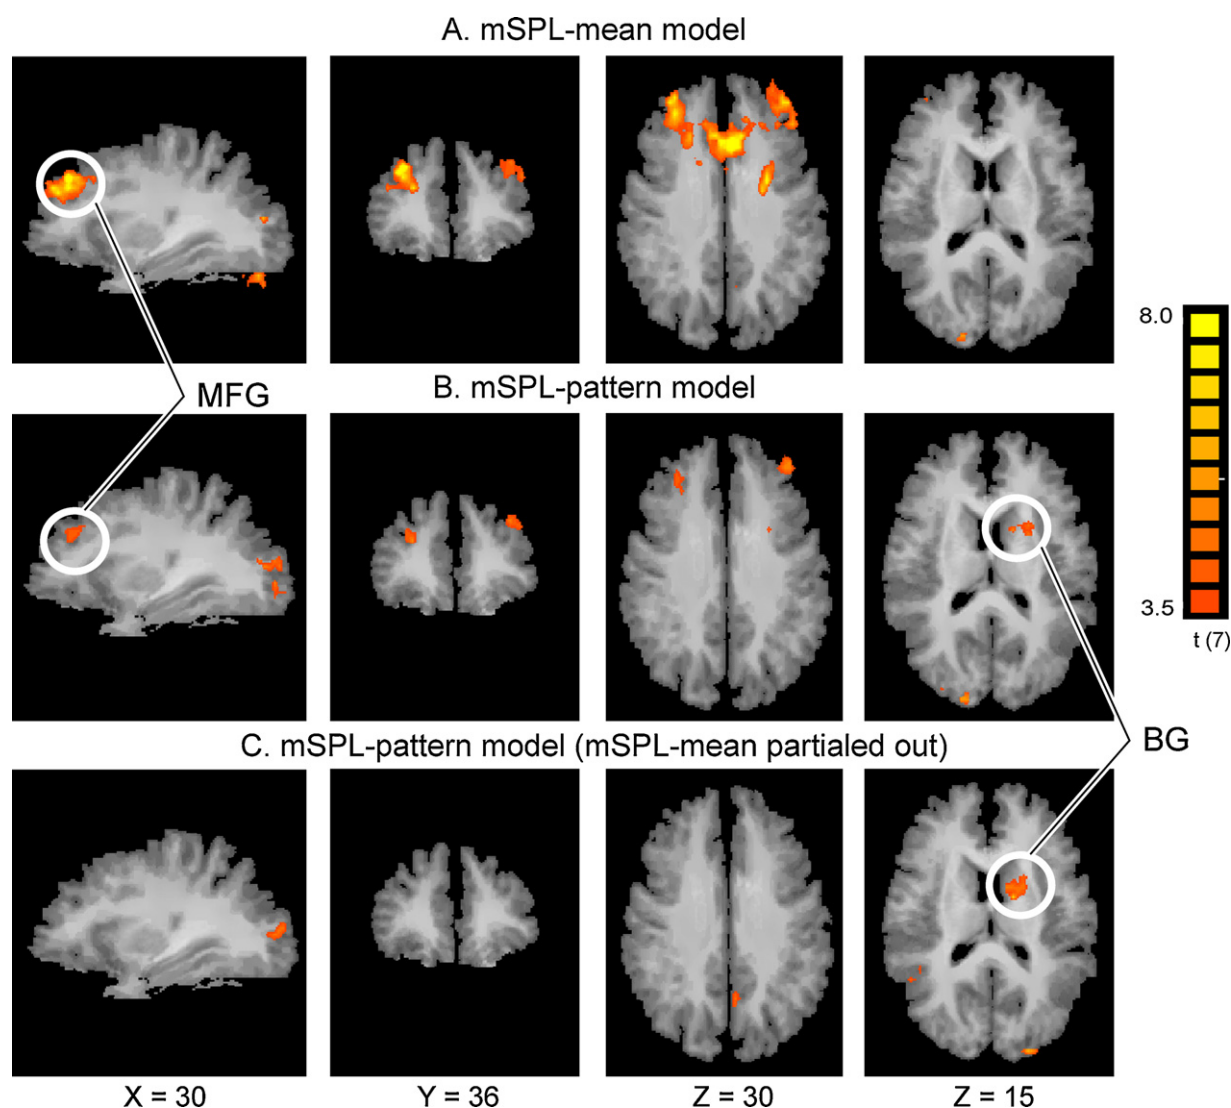


Fig. 3. Comparison of two functional connectivity analyses. (A) Conventional functional connectivity analysis results using the mean BOLD signal time course in mSPL as the seed (mSPL-mean model). The correlated regions are quite extensive both in lateral and medial frontal lobes. (B) Correlation map from the mSPL MVPTC (mSPL-pattern model) revealing more focused clusters of activation and a previously undetected cluster in the basal ganglia (BG). (C) Correlation map of voxels uniquely correlated with the mSPL-pattern model, revealing clusters in the BG, similar to those in (B). The cluster size correction for multiple comparisons yielded the same family-wise Type I error rate of $p < 0.05$ for both sets of maps shown here. Note that the mSPL seed region and the surrounding parietal regions were masked out of both analyses. MFG, middle frontal gyrus.

Table 1
Peak Activation Coordinates from the Group-level Analyses.

	Side	Talairach coordinates ^a		
		Peak X	Peak Y	Peak Z
<i>mSPL mean functional connectivity</i>				
Middle frontal gyrus (MFG)	R	29	37	30
Superior frontal gyrus (SFG)	L	-32	34	35
Medial frontal gyrus		1	25	31
Cuneus	R	26	-79	16
<i>mSPL-pattern functional connectivity</i>				
Middle frontal gyrus (MFG)	R	28	35	29
Superior frontal gyrus (SFG)	L	-33	34	35
Caudate	L	-19	9	17
Cuneus	R	23	-79	18

^a Voxelwise nominal $p = 0.009$, $t(7) = 3.5$, corrected alpha = 0.05.

functional connectivity model, that analysis also revealed much more extensive (i.e., less selective) activations in frontal lobes than did the mSPL-pattern model. In addition, no BG activation was

observed in the mSPL-mean model even with a liberal threshold of $t(7) = 2.3$, voxelwise $p < 0.06$, uncorrected. This suggests that the mSPL-pattern model was more sensitive to BG involvement, and revealed more specific (focal) prefrontal activity than the SPL-mean model (cf. Fig. 3A and B).

Furthermore, in order to directly contrast the explanatory power of the mSPL-pattern model to the mSPL-mean model, we conducted a step-wise regression analysis procedure (to avoid issues of collinearity between the two time courses) using both mSPL-mean and mSPL-pattern time courses. Specifically, our first-stage GLM included mSPL-mean as a regressor (same procedure as above), but instead of obtaining the activation map, we obtained the residuals. In the second stage, the mSPL-pattern vector was entered in a GLM as a regressor of the mean-model residuals, yielding a new map as shown in Fig. 3C. This analysis revealed significant correlation in the BG, demonstrating that mSPL-pattern model uniquely accounted for activation in the basal ganglia, above and beyond that accounted for by the mSPL-mean model (cf. Fig. 3B and C).

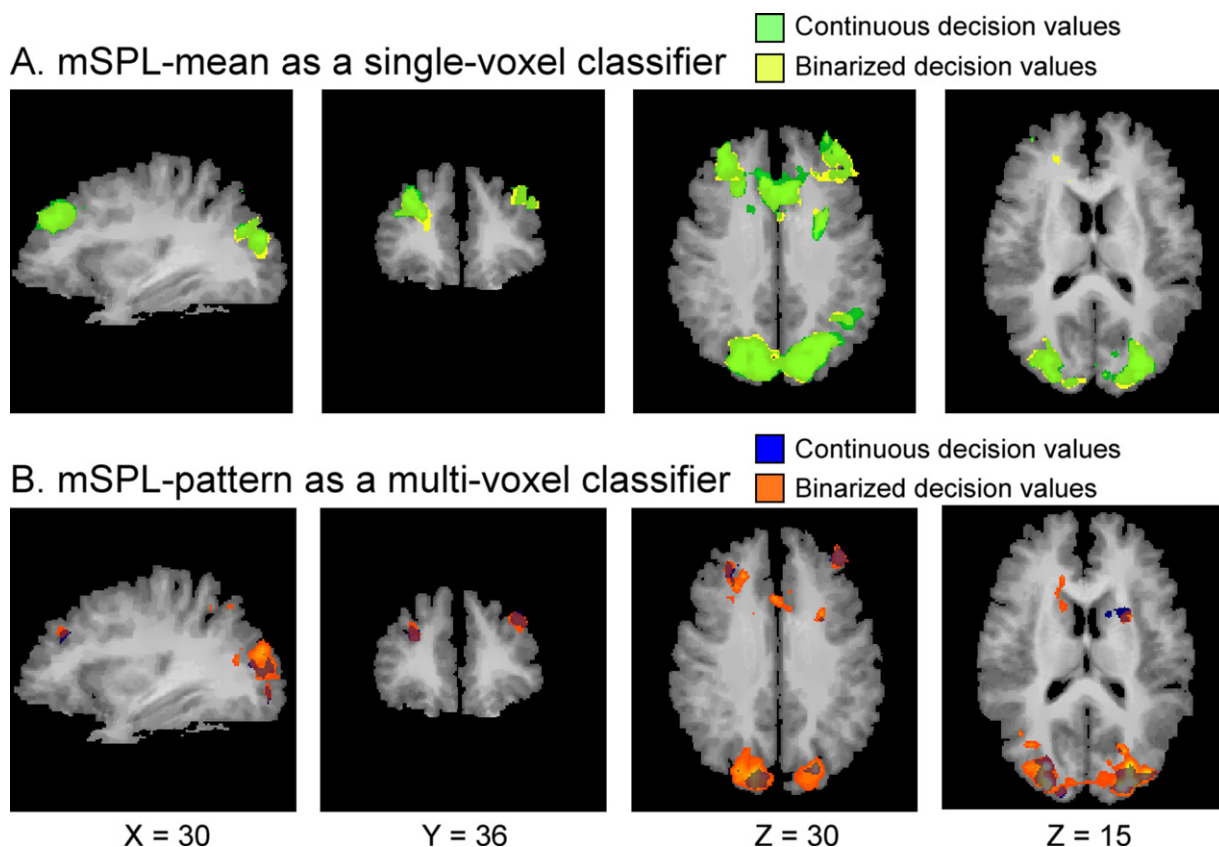


Fig. 4. Comparison of connectivity analysis using continuous vs. binarized decision values. The time course regressor was obtained from: (A) training SVM with mSPL-mean as a single-voxel classifier and creating continuous decision values (green) or binarized decision values (yellow); and from (B) training SVM with mSPL-pattern as multi-voxel classifier and creating continuous decision values (blue) or binarized values (orange). Note that the mSPL seed region and the surrounding parietal regions were not masked out here to ensure that the correlation maps were sensible.

Finally, it is important to ensure that the observed basal ganglia (BG) correlation was due to the use of multivariate pattern information, rather than some other aspect of the SVM training procedure. We thus conducted the following auxiliary analyses to exclude this possibility. First, we applied the SVM procedure to the SPL-mean as a single-feature classifier. We used both the raw/continuous decision values as well as the binarized values (Fig. 4A). No significant BG correlations were revealed, suggesting that the pattern contributed uniquely to the BG correlation. We also conducted the binarized MVPTC analysis as a comparison to the original MVPTC analysis with the continuous decision values (Fig. 4B). Both of these analyses revealed BG activation, suggesting that it was the pattern that provided the sensitivity.

4. Discussion

In this study, we replicated the finding of transient activity in mSPL associated with voluntary shifts of spatial attention (Fig. 2A; e.g., Kelley et al., 2008; Yantis et al., 2002). We then implemented a novel multivoxel pattern-based connectivity analysis method that revealed regions of the brain whose activity was correlated with fluctuations in multivoxel patterns within mSPL. These regions included prefrontal cortex and the striatum of the basal ganglia, providing converging evidence for an expanded parietal-frontal-striatal network for the shifting of spatial attention (e.g., Gitelman et al., 1999; Shulman et al., 2009).

We compared the results of the MVPTC correlation analysis (mSPL-pattern model) with a conventional functional connectivity analysis (mSPL-mean model) in which the time course of the mean BOLD signal in mSPL was used as a regressor. The mSPL-mean

model revealed extensive activity in the frontal regions (Fig. 3A). Studies have shown that the precuneus (of which mSPL is the dorsal-most part) is associated with a variety of cognitive functions (e.g., Cavanna & Trimble, 2006) in addition to the control of visuospatial attention shifts. Thus, some of the variance in the mean mSPL signal over time is likely to be driven by functions other than the control of attention shifts. In addition, there are intrinsic fluctuations in widespread cortical networks when participants are at rest (Aguirre, Zarahn, & D'Esposito, 1998; Macey, Macey, Kumar, & Harper, 2004) that could also account for connectivity in the mSPL-mean model.

In contrast, the mSPL MVPTC reflects changes in the mSPL multivoxel pattern that is associated specifically with the state of visuospatial attention; it was correlated with only a subset of the regions identified by the mSPL-mean model. Furthermore, this analysis uniquely revealed functional connectivity between the striatum and mSPL (see Table 1, cf., Fig. 3A and B). It is noteworthy that these results rule out the possibility that the regions identified with the mSPL-pattern model were simply the regions found to be most statistically robust (i.e., those with the highest t values) in the mSPL-mean model. The activity in the caudate detected by the mSPL-pattern model was not detected by the mSPL-mean model even with a more liberal threshold (voxel $p < 0.06$, uncorrected). We suggest that the MVPTC method reveals more process-specific functional connectivity, as the mSPL MVPTC specifically indexes shifting and sustained states of visuospatial attention, whereas the mean BOLD signal over the course of a run reflects a combination of multiple task-relevant, and possibly task-irrelevant, processes throughout the entire experiment.

The MVPTC analysis appears to be more sensitive than the HRF-based univariate GLM (cf. Figs. 3A and 2A); it revealed frontostriatal brain regions associated with attention shifting that were not revealed by the HRF-based univariate GLM. This difference could not be explained trivially by the choice of statistical threshold. An important feature of the mSPL-pattern model is that it generalizes to time points in the experiment beyond those used for training the classifier (i.e., during cued shifts of attention). For example, positive deflections in the MVPTC could reflect not only cued shifts of attention, but also spontaneous shifts of attention that are likely to occur from time to time (e.g., Christoff, Gordon, Smallwood, Smith, & Schooler, 2009; Weissman, Roberts, Visscher, & Woldorff, 2006). Furthermore, even for cued shifts of attention, the MVPTC may reflect meaningful variability in the amplitude and timing of each shift following the cue onset. In contrast, the conventional HRF-based univariate GLM assumes a consistent hemodynamic response function for each modeled attention-shift cue, and does not account for potential spontaneous shifts of attention.

The results of the mSPL MVPTC correlation analysis corroborate and extend previous findings in the literature suggesting a role for frontostriatal circuits in shifting between behaviorally relevant locations or objects (e.g., Cools et al., 2006, 2004; Shulman et al., 2009). Studies of patients with Parkinson's disease (which is associated with progressive neurodegeneration of the BG) exhibit deficits in set shifting (Ravizza & Ivry, 2001) as well as deficits in goal-directed, but not stimulus-driven, spatial attention (Grande et al., 2006). It has also been suggested that this circuit plays a role in selection and updating within working memory (O'Reilly & Frank, 2006). In fact, all shifts of goal-directed attention might be accompanied by an update in working memory, as suggested by O'Reilly and Frank's (2006) model.

The method described here is related to the hybrid SVM–GLM approach proposed by LaConte, Strother, Cherkassky, Anderson, and Hu (2005) and Wang (2009). Wang (2009) suggested that the hybrid SVM–GLM approach has better sensitivity than a conventional GLM for detecting task-relevant activations. Our method is also a hybrid SVM–GLM approach; however, the MVPTC analysis introduces two crucial extensions. First, while LaConte et al. (2005) and Wang (2009) both created a similar “pattern time course,” they did not use it for functional connectivity. Second, we used a rapid event-related design, whereas the previous implementations employed block designs (LaConte et al., 2005; Wang, 2009). Thus, we demonstrate that MVPTC can be used to make inferences about time points (TRs) that are not used in the training of the classifier model.

Our application avoids circular inferences (Kriegeskorte et al., 2009; Vul et al., 2009) by limiting the univariate GLM analysis to data from voxels that are spatially independent of the functionally defined ROI used to create the multivoxel-pattern model. Another recently proposed method for assessing functional connectivity uses multivoxel patterns to compute mutual information between two or more regions of interest (Chai, Walther, Beck, & Fei-Fei, 2009); this is different both from the MVPTC approach proposed here and from the methods of LaConte et al. (2005) and Wang (2009), but points to a number of potential novel uses of MVPA beyond simple decoding.

The present study introduces a new way to use multivoxel pattern analysis that exploits evolving patterns of activity during task performance. Multivoxel pattern analysis is first used to distinguish patterns of activity associated with cognitive processes of interest, and the resulting classifier is then used to compute a multivoxel pattern time course that reflects fluctuations in cognitive state over time. Although here we modeled transient attention-related activity, there are a wide range of potential applications of this method. These include modeling sustained attention to dif-

ferent spatial locations or objects and modeling working memory operations for different stimuli or task sets.

A major advantage of this method is that it could provide a way to track fluctuating brain states during tasks in which the timing and effectiveness of a given cognitive operation (i.e., how rapid and accurate was a given shift of attention?) may not be closely tied to external experimental events. For example, one could train a multivoxel pattern model to discriminate between patterns of activity associated with high task engagement vs. patterns of activity associated with low task engagement (e.g., during mind-wandering, Christoff et al., 2009; or attentional lapses, Weissman et al., 2006), and then apply the estimated weights to data from an independent data set to predict fluctuations in behavioral performance. These and other potential applications of the MVPTC may provide new insights into how ongoing changes in brain state give rise to fluctuations in cognition and behavior.

Acknowledgments

This work was supported by National Institutes of Health Grant R01-DA013165 to S.Y. and postdoctoral National Research Service Awards from the NEI (T32EY07143) and the NIA (AG027668-01) to L.G. We thank Dr. Amy Shelton for helpful discussions and James Gao for research assistance.

References

- Aguirre, G. K., Zarahn, E., & D'Esposito, M. (1998). The inferential impact of global signal covariates in functional neuroimaging analyses. *Neuroimage*, 8, 302–306.
- Biswal, B., Yetkin, F. Z., Haughton, V. M., & Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine*, 34, 537–541.
- Boynton, G. M., Engel, S. A., Glover, G. H., & Heeger, D. J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *Journal of Neuroscience*, 16, 4207–4221.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433–436.
- Cavanna, A. E., & Trimble, M. R. (2006). The precuneus: A review of its functional anatomy and behavioural correlates. *Brain*, 129, 564–583.
- Chai, B., Walther, D. B., Beck, D. M., & Fei-Fei, L. (2009). *Exploring functional connectivity of the human brain using multivariate information analysis*. Neural Information Processing System (NIPS).
- Chang, C.-C., & Lin, C.-J. (2001). *LIBSVM: A library for support vector machines*. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>
- Chiu, Y. C., Esterman, M., Han, Y., Rosen, H., & Yantis, S. (2011). Decoding task-based attentional modulation during face categorization. *Journal of Cognitive Neuroscience*, 23, 1198–1204.
- Chiu, Y. C., & Yantis, S. (2009). A domain-independent source of cognitive control for task sets: Shifting spatial attention and switching categorization rules. *Journal of Neuroscience*, 29, 3930–3938.
- Christoff, K., Gordon, A. M., Smallwood, J., Smith, R., & Schooler, J. W. (2009). Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. *Proceedings of the National Academy of Sciences of the United States of America*, 106, 8719–8724.
- Cools, R., Altamirano, L., & D'Esposito, M. (2006). Reversal learning in Parkinson's disease depends on medication status and outcome valence. *Neuropsychologia*, 44, 1663–1673.
- Cools, R., Clark, L., & Robbins, T. W. (2004). Differential responses in human striatum and prefrontal cortex to changes in object and rule relevance. *Journal of Neuroscience*, 24, 1129–1135.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3, 201–215.
- Diekhof, E. K., Falkai, P., & Gruber, O. (2009). Functional interactions guiding adaptive processing of behavioral significance. *Human Brain Mapping*, 30, 3325–3331.
- Duann, J. R., Ide, J. S., Luo, X., & Li, C. S. (2009). Functional connectivity delineates distinct roles of the inferior frontal cortex and presupplementary motor area in stop signal inhibition. *Journal of Neuroscience*, 29, 10171–10179.
- Esterman, M., Chiu, Y. C., Tamber-Rosenau, B. J., & Yantis, S. (2009). Decoding cognitive control in human parietal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 106, 17974–17979.
- Forman, S. D., Cohen, J. D., Fitzgerald, M., Eddy, W. F., Mintun, M. A., & Noll, D. C. (1995). Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): Use of a cluster-size threshold. *Magnetic Resonance in Medicine*, 33(5), 636–647.
- Fox, M. D., Corbetta, M., Snyder, A. Z., Vincent, J. L., & Raichle, M. E. (2006). Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. *Proceedings of the National Academy of Sciences of the United States of America*, 103, 10046–10051.

- Fox, M. D., Zhang, D., Snyder, A. Z., & Raichle, M. E. (2009). The global signal and observed anticorrelated resting state brain networks. *Journal of Neurophysiology*, *101*, 3270–3283.
- Friston, K. J., Buechel, C., Fink, G. R., Morris, J., Rolls, E., & Dolan, R. J. (1997). Psychophysiological and modulatory interactions in neuroimaging. *Neuroimage*, *6*, 218–229.
- Friston, K. J., Frith, C. D., Liddle, P. F., & Frackowiak, R. S. (1993). Functional connectivity: The principal-component analysis of large (PET) data sets. *Journal of Cerebral Blood Flow and Metabolism*, *13*, 5–14.
- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J. B., Frith, C., & Frackowiak, R. S. J. (1995). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, *2*, 189–210.
- Gitelman, D. R., Nobre, A. C., Parrish, T. B., LaBar, K. S., Kim, Y. H., Meyer, J. R., et al. (1999). A large-scale distributed network for covert spatial attention: Further anatomical delineation based on stringent behavioural and cognitive controls. *Brain*, *122*, 1093–1106.
- Grande, L. J., Crosson, B., Heilman, K. M., Bauer, R. M., Kilduff, P., & McGlinchey, R. E. (2006). Visual selective attention in Parkinson's disease: Dissociation of exogenous and endogenous inhibition. *Neuropsychology*, *20*, 370–382.
- Greenberg, A. S., Esterman, M., Wilson, D., Serences, J. T., & Yantis, S. (2010). Control of spatial and feature-based attention in frontoparietal cortex. *Journal of Neuroscience*, *30*, 14330–14339.
- He, B. J., Shulman, G. L., Snyder, A. Z., & Corbetta, M. (2007). The role of impaired neuronal communication in neurological disorders. *Current Opinion in Neurology*, *20*, 655–660.
- Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, *8*, 679–685.
- Kelley, T. A., Serences, J. T., Giesbrecht, B., & Yantis, S. (2008). Cortical mechanisms for shifting and holding visuospatial attention. *Cerebral Cortex*, *18*, 114–125.
- Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences of the United States of America*, *103*, 3863–3868.
- Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S., & Baker, C. I. (2009). Circular analysis in systems neuroscience: The dangers of double dipping. *Nature Neuroscience*, *12*, 535–540.
- LaConte, S., Strother, S., Cherkassky, V., Anderson, J., & Hu, X. (2005). Support vector machines for temporal classification of block design fMRI data. *Neuroimage*, *26*, 317–329.
- Leber, A. B., Turk-Browne, N. B., & Chun, M. M. (2008). Neural predictors of moment-to-moment fluctuations in cognitive flexibility. *Proceedings of the National Academy of Sciences of the United States of America*, *105*, 13592–13597.
- Liu, T., Slotnick, S. D., Serences, J. T., & Yantis, S. (2003). Cortical mechanisms of feature-based attentional control. *Cerebral Cortex*, *13*, 1334–1343.
- Macey, P. M., Macey, K. E., Kumar, R., & Harper, R. M. (2004). A method for the removal of global effects from fMRI time series. *Neuroimage*, *22*, 360–366.
- Norman, K. A., Polyn, S. M., Detre, G. L., & Haxby, J. V. (2006). Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Science*, *10*, 424–430.
- O'Reilly, R. C., & Frank, M. J. (2006). Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation*, *18*, 283–328.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Polyn, S. M., Natu, V. S., Cohen, J. D., & Norman, K. A. (2005). Category-specific cortical activity precedes retrieval during memory search. *Science*, *310*, 1963–1966.
- Ravizza, S. M., & Ivry, R. B. (2001). Comparison of the basal ganglia and cerebellum in shifting attention. *Journal of Cognitive Neuroscience*, *13*, 285–297.
- Rissman, J., Gazzaley, A., & D'Esposito, M. (2004). Measuring functional connectivity during distinct stages of a cognitive task. *Neuroimage*, *23*, 752–763.
- Saygin, A. P., & Sereno, M. I. (2008). Retinotopy and attention in human occipital, temporal, parietal, and frontal cortex. *Cerebral Cortex*, *18*, 2158–2168.
- Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological Science*, *20*, 207–214.
- Serences, J. T., Schwarzbach, J., Courtney, S. M., Golay, X., & Yantis, S. (2004). Control of object-based attention in human cortex. *Cerebral Cortex*, *14*, 1346–1357.
- Serences, J. T., & Yantis, S. (2007). Representation of attentional priority in human occipital, parietal, and frontal cortex. *Cerebral Cortex*, *17*, 284–293.
- Shomstein, S., & Yantis, S. (2006). Parietal cortex mediates voluntary control of spatial and nonspatial auditory attention. *Journal of Neuroscience*, *26*, 435–439.
- Shulman, G. L., Astafiev, S. V., Franke, D., Pope, D. L., Snyder, A. Z., McAvoy, M. P., et al. (2009). Interaction of stimulus-driven reorienting and expectation in ventral and dorsal frontoparietal and basal ganglia-cortical networks. *Journal of Neuroscience*, *29*, 4392–4407.
- Silver, M. A., Ress, D., & Heeger, D. J. (2005). Topographic maps of visual spatial attention in human parietal cortex. *Journal of Neurophysiology*, *94*, 1358–1371.
- Smith, A. P., Stephan, K. E., Rugg, M. D., & Dolan, R. J. (2006). Task and content modulate amygdala–hippocampal connectivity in emotional retrieval. *Neuron*, *49*, 631–638.
- Stephan, K. E., Marshall, J. C., Penny, W. D., Friston, K. J., & Fink, G. R. (2007). Inter-hemispheric integration of visual processing during task-driven lateralization. *Journal of Neuroscience*, *27*, 3512–3522.
- Tamber-Rosenau, B. J., Esterman, M., Chiu, Y. C., & Yantis, S. (2011). Cortical mechanisms of cognitive control for shifting attention in vision and working memory. *Journal of Cognitive Neuroscience*, www.mitpressjournals.org/doi/pdf/10.1162/jocn.2011.21608
- Thiebaut de Schotten, M., Urbanski, M., Duffau, H., Volle, E., Lévy, R., Dubois, B., et al. (2005). Direct evidence for a parietal–frontal pathway subserving spatial awareness in humans. *Science*, *309*, 2226–2228.
- Vandenberghe, R., Gitelman, D. R., Parrish, T. B., & Mesulam, M. M. (2001). Functional specificity of superior parietal mediation of spatial shifting. *Neuroimage*, *14*, 661–673.
- Vul, E., Harris, C., Winkielman, P., & Pashler, H. (2009). Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition. *Perspectives on Psychological Science*, *4*, 274–290.
- Wang, Z. (2009). A hybrid SVM–GLM approach for fMRI data analysis. *Neuroimage*, *46*, 608–615.
- Weissman, D. H., Roberts, K. C., Visscher, K. M., & Woldorff, M. G. (2006). The neural bases of momentary lapses in attention. *Nature Neuroscience*, *9*, 971–978.
- Yantis, S., Schwarzbach, J., Serences, J. T., Carlson, R. L., Steinmetz, M. A., Pekar, J. J., et al. (2002). Transient neural activity in human parietal cortex during spatial attention shifts. *Nature Neuroscience*, *5*, 995–1002.
- Zarahn, E., Aguirre, G. K., & D'Esposito, M. (1997). Empirical analyses of BOLD fMRI statistics. I. Spatially unsmoothed data collected under null-hypothesis conditions. *Neuroimage*, *5*, 179–197.