Tracking the will to attend: Cortical activity indexes self-generated, voluntary shifts of attention

Leon Gmeindl[1] · Yu-Chin Chiu[1,2] · Michael S. Esterman[1] · Adam S. Greenberg[1] · Susan M. Courtney[1] · Steven Yantis[1]

Abstract The neural substrates of volition have long tantalized philosophers and scientists. Over the past few decades, researchers have employed increasingly sophisticated technology to investigate this issue, but many studies have been limited considerably by their reliance on intrusive experimental procedures (e.g., abrupt instructional cues), measures of brain activity contaminated by overt behavior, or introspective self-report techniques of questionable validity. Here, we used multivoxel pattern time-course analysis of functional magnetic resonance imaging data to index voluntary, covert perceptual acts—shifts of visuospatial attention—in the absence of instructional cues, overt behavioral indices, and self-report. We found that these self-generated, voluntary attention shifts were time-locked to activity in the medial superior parietal lobule, supporting the hypothesis that this brain region is engaged in voluntary attentional reconfiguration. Self-generated attention shifts were also time-locked to activity in the basal ganglia, a novel finding that motivates further research into the role of the basal ganglia in acts of volition. Remarkably, *prior* to self-generated shifts of attention, we observed early and selective increases in the activation of medial frontal (dorsal anterior cingulate) and lateral prefrontal (right middle frontal gyrus) cortex—activity that likely reflects processing related to the intention or preparation to reorient attention. These findings, which extend recent evidence on freely chosen motor movements, suggest that dorsal anterior cingulate and lateral prefrontal cortices play key roles in both overt and covert acts of volition, and may constitute core components of a brain network underlying the will to attend.

Keywords Neural attention mechanisms · Cognitive neuroscience · Imaging

Humans have a highly developed ability to attend voluntarily to some stimuli while ignoring others—an ability that supports a spectrum of behavior, from searching for objects in a cluttered environment to multitasking and problem solving. Impaired performance arising from diminished attentional control is associated with aging, neurological diseases and disorders (e.g., attention deficit hyperactivity disorder), and substance abuse. Yet, fundamental questions about how the brain implements voluntary control of attention persist.

Functional magnetic resonance imaging (fMRI) studies have indicated that the medial superior parietal lobule (mSPL), a component of the dorsal attention network (Corbetta & Shulman, 2002), is activated during shifts of attention between spatial locations (Greenberg, Esterman, Wilson, Serences, & Yantis, 2010; Kelley, Serences, Giesbrecht, & Yantis, 2008; Vandenberghe, Gitelman, Parrish, & Mesulam, 2001; Yantis et al., 2002), features (Greenberg, Esterman, Wilson, Serences, & Yantis, 2010; Liu, Slotnick, Serences, & Yantis, 2003), objects (Serences, Schwarzbach, Courtney, Golay, & Yantis, 2004), sensory modalities (Shomstein & Yantis, 2004), task sets (Chiu & Yantis, 2009), and working memory representations (Tamber-Rosenau, Esterman, Chiu, & Yantis, 2011). However, in these and other studies (Hopfinger, Camblin, & Parks, 2010; Taylor,

Leon Gmeindl and Yu-Chin Chiu contributed equally to this work.

✉ Leon Gmeindl
gmeindl@jhu.edu

✉ Yu-Chin Chiu
chiu.yuchin@duke.edu

[1] Psychological and Brain Sciences, Johns Hopkins University, 3400 N. Charles Street, Ames Hall, Baltimore, MD 21218-2680, USA

[2] Center for Cognitive Neuroscience, Duke University, LSRC Box 90999, Durham, NC 27705, USA

Rushworth, & Nobre, 2008), participants were presented with instructional cues indicating when to shift attention, or they were required to report when they had voluntarily shifted attention by making a motor response. These methods introduce at least two substantial challenges regarding analysis and interpretation: First, external cues and behavioral responses evoke their own patterns of neural activity that somehow must be dissociated from the activity associated with the intention or preparation to perform an action. Second, initiating ostensibly voluntary actions in response to external cues is not purely voluntary; rather, it is partly stimulus-driven. It is therefore unclear whether or under what circumstances the observed activity in mSPL reflects a purely voluntary orienting of attention, or instead involves the anticipation and subsequent processing of an external instruction (e.g., processes involved in cue interpretation and stimulus–response mapping).

Here we introduce a novel method to track attention in the absence of both external cues and overt behavior, allowing us to isolate the neural mechanisms engaged in self-generated, voluntary shifts of attention. Participants were instructed to fix their gaze on a central fixation stimulus while covertly attending alternately to one of two locations throughout a continuous multistream rapid serial visual presentation (RSVP) task (Fig. 1). In the "uncued" condition, participants were further instructed to shift attention voluntarily (without moving their eyes) roughly three or four times per minute; they were not required to execute these shifts at any particular times during the task, and no instructional cues were used to evoke attention shifts. Furthermore, participants did not indicate explicitly (e.g., by motor response or retrospective report; Hopfinger, Camblin, & Parks, 2010; Lau, Rogers, & Passingham, 2006; Libet, Gleason, Wright, & Pearl, 1983) their intention or decision to shift attention, the time at which the shift occurred, or the location to which they shifted attention. Instead, we exploited the fact that attention to different spatial locations evokes systematically distinct patterns of brain activity in visual cortex. This activity was used to track the focus of attention over time.

Specifically, we first trained a multivoxel pattern analysis (MVPA) classifier to distinguish the visual cortex activity associated with covert attention to left versus right locations by using the fMRI data from a separate "cued" condition, in which participants occasionally were cued to attend to the left or the right location, thus eliciting cued shifts of attention on a subset of trials (Fig. 1). In the cued condition, participants were instructed to shift attention only when the cues were presented. The classifier trained on data from the cued condition was then was applied to the fMRI data from the uncued condition, thereby tracking the focus of attention on a moment-by-moment basis. At each time point, the participant's attention was classified as being oriented to the left or to the right location. We then demarcated time points when the

activation patterns indicated a transition from a leftward to a rightward focus of attention, or vice versa. These temporal markers of attention shifts were then employed in subsequent analyses to identify the neural mechanisms engaged for self-generated shifts of attention.

Furthermore, we analyzed the time course of activity in some brain regions that were engaged for both uncued and cued shifts of attention, to test whether there was an earlier increase in activity in these regions prior to self-generated shifts of attention—an increase that should reflect the preparation or intention (i.e., the will) to shift attention.

Method

Ethics statement

All experimental procedures were approved by the Johns Hopkins Medical Institutional Review Board. All participants passed an fMRI safety screening prior to the scan and provided written informed consent.

Participants

Twelve neurologically intact, right-handed, healthy adults (seven females, five males; mean age = 26, SE = 1.9 years) were recruited from the Johns Hopkins University community. All participants had normal or corrected-to-normal vision. Each participant completed one behavioral training session and one 2-h scanning session, on separate days. They were paid $10/h for the behavioral training session and $25/h for the scanning session.

Apparatus

In the training session, visual stimuli were presented on an 18-in. CRT monitor located 79 cm in front of a chinrest, used to equate visual angles across participants, and buttonpress responses were made on a computer keyboard. In the fMRI session, the 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. Each participant was fitted with a custom-molded dental impression block clamped to the head coil cage, to minimize head motion; buttonpress responses were made on a custom-built MR-compatible response box. Stimulus presentation and behavioral data collection were controlled by custom MATLAB (The MathWorks, Inc.) code using the Psychophysics Toolbox (Brainard, 1997). Eye position was monitored with a closed-circuit video system during the practice session, and with a custom MR-compatible infrared camera (MRA, Inc.) and ViewPoint 2.8.3 eyetracking software (Arrington Research, Inc.) during the fMRI session.
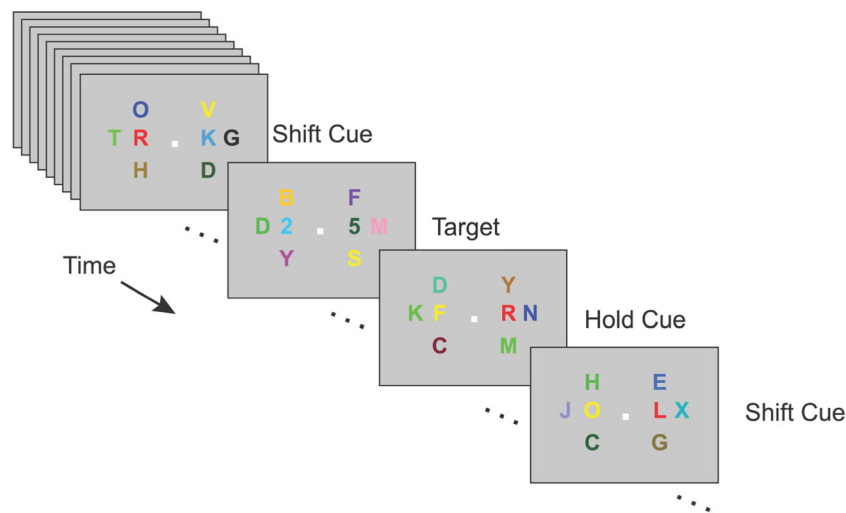
Fig. 1 Task. Two task-relevant streams of characters (immediately to the left and right of a central fixation) were flanked by task-irrelevant streams. In the *cued-attention* condition (illustrated here), uniquely red stimuli (L and R, corresponding to left and right, respectively) cued participants either to shift attention to the other stream (e.g., an R cue appearing in the left stream) or to hold attention on the currently relevant stream (e.g., an R cue appearing in the right stream). Participants indicated the identity of targets (digits) that appeared within the currently attended stream. In the *uncued-attention* condition, the cues were omitted; participants were to shift attention between the relevant streams occasionally and at will. In both conditions, digits occurred only rarely, and they appeared simultaneously in both the left and right streams

Stimuli and procedures

Participants were instructed to fix their gaze on a white central fixation dot 0.2° in diameter while performing a multistream RSVP task (Fig. 1). Task-relevant alphanumeric RSVP streams were located 3.5° to the left and right of the fixation dot along the horizontal meridian. Each of these two relevant streams was flanked 3.3° (center to center) above, below, and laterally by three irrelevant distractor streams in order to maximize the demand for selective attention. The alphanumeric characters subtended 1.4° in height and 1.0° in width and were presented in fixed-width Monaco font (letters in uppercase). Participants were instructed to make four-alternative buttonpress responses to infrequent target digits embedded within the task-relevant RSVP streams. Simultaneously, a digit from 2 to 5 was presented in one stream and a different digit from 2 to 5 was presented in the other stream. Participants pressed the right index-, middle-, ring-, or little-finger button to indicate the identity of the digit (2, 3, 4, or 5, respectively) presented in the currently attended RSVP stream. The filler (nontarget) items consisted of the letters A through Z, except for L and R (see below). All visual stimuli were presented on a gray background, and each target and filler item was rendered in one of eight randomly chosen colors (excluding red), with the constraint that every item within the same RSVP frame was rendered in a different color. The stimulus duration (i.e., RSVP frame duration) was 133.3 ms, and targets were infrequent, appearing on average two or three times per minute. The stimuli requiring motor responses were rare and included only to ensure that the participants remained vigilant.

Each fMRI scanning run consisted of one trial block of the cued condition and one trial block of the uncued condition (see below). Block order was counterbalanced across runs and across participants, and printed instructions ("red cues" or "self-paced") indicating the relevant condition were presented immediately prior to each block. Each run lasted 410.2 s (189.1 s per block), including an initial 12-s fixation period, as well as a 20-s fixation period inserted between the two blocks. Each participant completed ten runs in the fMRI session, conducted at the F. M. Kirby Research Center for Functional Brain Imaging in Baltimore, Maryland.

In the cued condition, participants occasionally were instructed either to shift attention to the other task-relevant RSVP stream or to maintain attention ("hold") on the current stream. Shift and hold instructions were conveyed by the letters L and R rendered in bright red, which appeared unpredictably within the currently relevant stream. When an R appeared in the left stream, participants were to shift attention to the right stream; in contrast, when an L appeared in the left stream, participants were to maintain attention to the left stream—L and R signaled the reverse instructions when presented in the right stream. Only the shift and hold cues were rendered in bright red, so that participants could easily discriminate cues from the other stimuli; this discriminability was verified during the practice session. The onset asynchrony between critical events (i.e., the presentation of a cue or a target) in the cued condition varied randomly among 5.067, 6.000, 7.067, 8.000, 9.067, or 10.000 s. Approximately half of the cues were shift cues, and approximately half were hold cues. A variable number of hold cues were presented between

successive shift cues (resulting in unpredictable cue sequences), and, of particular importance, the mean onset asynchrony between shift cues was 19.8 s (observed asynchrony range = 5.067–44.267 s, between-participants SD = 1.5 s). Targets were always separated from cues by a minimum onset asynchrony of 5.067 s.

In the uncued condition, the shift and hold cues were omitted (replaced by filler letters). Instead, participants were instructed to shift attention voluntarily from one task-relevant stream to the other a few (roughly three or four) times each minute, and to respond to targets appearing within the currently attended (relevant) stream just as they had responded in the cued condition. Participants were not required to shift attention at any particular times in the uncued condition. A previous study (Gmeindl, Gao, Yantis, & Courtney, 2008) employing a very similar design, but one in which buttonpress responses were used to verify the accurate timing of the indexed attention shifts, indicated that with these instructions participants shifted attention between the left and right RSVP streams every 21.1 s on average (SE = 2.3 s).

Prior to the fMRI session, each participant completed a training session in our laboratory at Johns Hopkins University. Throughout the training session, the RSVP stimulus frame duration was incrementally decreased from 400 to 133.3 ms, a rate at which participants were able to maintain an accuracy in the cued condition of at least 80 % correct across two successive blocks. In the practice session, accuracy feedback was provided at the end of each block. In the fMRI session, the RSVP stimulus frame duration was fixed at 133.3 ms and accuracy feedback was omitted.

Imaging procedures

Data acquisition Functional MRI data were acquired with a Philips Intera 3-T scanner and an eight-channel SENSE head coil (MRI Devices). High-resolution, whole-brain anatomical volumes were acquired with an MPRAGE T1-weighted sequence yielding 200 1-mm coronal slices (1 × 1 mm in-plane resolution, matrix = 256 × 256, TE = 3.7 ms, TR = 8.1 ms, flip angle = 8°). Whole-brain functional volumes were acquired with a T2*-weighted echoplanar imaging sequence yielding 30 2.5-mm axial slices (1-mm gap, 2.5 × 2.5 mm in-plane resolution, matrix = 76 × 76, TE = 30 ms, TR = 1.5 s, flip angle = 70°). Eight subsequently discarded volumes were collected at the beginning of each run to allow magnetization to reach a steady state prior to task presentation.

Imaging data preprocessing The functional MRI data were preprocessed using the BrainVoyager QX software, version 1.10 (Brain Innovation). The data from each run were corrected for slice-time acquisition and motion, and then temporally high-pass filtered (three cycles per run). To correct for between-run motion, each participant's functional volumes were all coregistered to his or her high-resolution anatomical volume. Voxels were resampled to 3 × 3 × 3 mm. No other spatial smoothing or normalization was performed. After preprocessing, the blood oxygenation level dependent (BOLD) time course was extracted from each voxel for each run. The BOLD amplitude at each time point (i.e., the functional volume, or TR) was transformed into a z score with respect to the mean and standard deviation of the voxel's time course for that run, and then entered into MATLAB for the MVPA (see below). To conduct the Cued-Shift > Cued-Hold contrast reported below, each participant's anatomical and functional volumes were transformed into Talairach space using a rigid-body transformation. Then, for each participant, a general linear model of the cued-attention data was formed. This model included regressors for shift cues, hold cues, and targets that were each created by convolving a single-gamma hemodynamic response function with Kronecker delta (stick) functions that marked the onsets of the corresponding events; head movements in the x, y, and z dimensions were included as regressors of no interest. A standard group-level analysis (one-sample t test) was then performed on the results of a Cued-Shift > Cued-Hold contrast, with statistical maps being corrected for multiple comparisons by applying a cluster-size threshold [voxel-wise p < .001, $t(11)$ = 3.5, corrected α = .05].

Multivoxel pattern analysis[1] For each participant, we first used data from only the cued condition to train an MVPA classifier (linear support vector machine LIBSVM; Chang & Lin, 2011; www.csie.ntu.edu.tw/~cjlin/libsvm) to distinguish between the patterns of activity associated with sustained attention to the left versus the right RSVP stream. These were relatively long epochs (up to 44.3 s) during which

---

[1] For researchers who want to use participants brain activity to track the focus of attention—but without implementing MVPA—it is worth noting that in a preliminary study we used an alternative method that likewise revealed activation in mSPL, right middle frontal gyrus, and dorsal anterior cingulate cortex related to uncued, self-generated shifts of attention. Briefly, for each participant we (1) extracted the BOLD time courses from regions that were found to be reliably activated (or deactivated), according to an Attend Left > Attend Right whole-brain contrast performed on the data from the cued condition; (2) at each time point (TR), computed the ratio of mean BOLD signals between (striate and extrastriate) regions that were reliably more active for Attend Left epochs versus those that were reliably more active for Attend Right epochs; (3) computed the first derivative of this ratio; (4) optimized a velocity (i.e., change in ratio over time) threshold used to demarcate cued attention shifts, such that it maximized hits (i.e., identifying high-velocity deflections in the BOLD ratio following presentation of the shift cues) and minimized false alarms; and (5) applied this optimized velocity threshold to the BOLD-ratio time course observed for these same regions in the *uncued* condition, thereby indexing self-generated shift-time points. These time points could then be used to create event-related average time courses as in the primary method reported above. Also, it may be noted that essentially the same methods, both MVPA-based and the alternative method above, could be used to track attention to some nonspatial stimulus features and objects (e.g., faces vs. scenes).

participants had been cued to attend to the RSVP stream on the left (Attend Left) or the right (Attend Right) of fixation. We trained the classifier on these Attend Left and Attend Right epochs using multivoxel patterns recorded from 7.5 s after the onset of the corresponding epoch (to account for the hemodynamic response lag) to 1.5 s after the offset of the corresponding epoch, with each constituent time point (i.e., TR) treated as a training sample. The MVPA was conducted first using all voxels within a whole-brain mask (ventricles excluded) created separately for each participant in the native anatomical space. The number of voxels in the mask ranged from 41,608 to 50,314 across participants ($M = 47,525$). A standard leave-one-run-out cross-validation procedure was used to evaluate classification accuracy. This initial whole-brain MVPA provided a set of weights (one per voxel), for each participant, that indicated how much information each voxel contributed to the correct classification. To select those voxels that were most informative, we ranked all of the voxels according to the absolute values of their weights and then repeated the cross-validation procedure with increasingly larger subsets (1, 50, 100, 500, 1,000, etc.) of the most informative voxels. Classification accuracy, averaged across participants, varied approximately as an inverted-U function of the number of voxels included in the MVPA, with a peak mean accuracy at 3,000 voxels. Therefore, we selected for each participant the 3,000 most informative voxels and trained a new, optimized classifier on the cued-condition data from these 3,000 voxels (no run left out).

This optimized classifier was then applied to the data from the uncued condition, resulting in a multivoxel pattern time course (MVPTC; Chiu, Esterman, Gmeindl, & Yantis, 2012; see also Greenberg et al., 2010) that indicated, for each time point, the degree to which the pattern of activity across the 3,000 voxels corresponded to the patterns associated with Attend Left versus Attend Right. The MVPTC was temporally smoothed by averaging the MVPTC value at each time point with the MVPTC values for the subsequent two time points, and then the MVPTC was binarized. The time points at which the classification reversed (i.e., shifting from left to right or vice versa) were demarcated as attention-shift points.

To verify that the MVPTC could be used to reliably index attention shifts, we used a leave-one-run-out cross-validation procedure in which, for each participant, we iteratively left out the data from one run of the cued condition (e.g., Run 1) and trained the classifier using the rest of the data from the cued condition (e.g., Runs 2–10). For each run left out, we then compared the onsets of attention shifts indexed by the MVPTC to the actual onsets of the shift cues. If the onset of an indexed attention shift fell within three TRs (i.e., 4.5 s, to account for hemodynamic response lag) of the onset of a shift cue, we considered this a hit. Across participants, the mean hit rate was 87.6 % (SE = 2.3 %). The false-alarm rate, as defined by the three-TR threshold, was comparatively low

($M = 34.3$ %, $SE = 5.1$ %) and may reflect that attention likely did fluctuate occasionally during performance (resulting, e.g., in missed targets). The hit rate was significantly higher than the false-alarm rate [$t(11) = 7.29$, $p < .001$], indicating that the MVPTC reliably indexed attention shifts across participants and across runs within each scanning session.

A priori regions of interest A recent study (Chiu et al., 2012) using a novel MVPTC analysis revealed that two cortical regions—right middle frontal gyrus (rMFG) and dorsal anterior cingulate cortex (dACC)—and one subcortical cluster in the basal ganglia (BG) demonstrated functional connectivity with the mSPL during cued shifts of attention. Furthermore, rMFG and dACC had also been implicated in our preliminary study (Gmeindl et al., 2008), in which participants engaged in uncued attention shifts during a similar task, but one in which buttonpress responses were used to verify the timing of the demarcated attention shifts. Of particular interest, that study indicated that rMFG and dACC were activated reliably more for self-generated than for cue-driven shifts. On the basis of these findings, we therefore included in the present study rMFG, dACC, and BG a priori regions of interest (ROIs; Table 1) that were functionally defined on the basis of the data from Chiu et al. (2012), and we tested for increased preparatory processing in these regions prior to self-generated shifts of attention.

Event-related average time-course analysis To test the directional hypothesis that self-generated attention shifts are associated with earlier rises in activity within the a priori ROIs (Table 1) than are cued attention shifts, we performed an event-related time-course analysis (Serences, 2004).

We first extracted time courses from the mSPL, rMFG, dACC, and BG regions (see Tables 1 and 2 and Fig. 2) by calculating the mean BOLD amplitude across all voxels within the region for each time point, covering 12 s centered on the uncued-shift time point identified by the MVPA classifier (i.e., from 6.0 s before to 6.0 s after the uncued-shift time point).

Table 1 A priori regions of interest

Talairach Coordinates				
X	Y	Z	# Voxels	Location of Peak Voxel
28	33	32	185	Right middle frontal gyrus
−1	21	27	62	Dorsal anterior cingulate
−16	8	15	87	Basal ganglia (region includes left caudate nucleus and putamen)

Table 2 Cued-Shift > Cued-Hold contrast

Talairach Coordinates

X	Y	Z	# Voxels	Location of Peak Voxel
−16	−71	51	281	Left precuneus
−31	−47	60	62	Left superior parietal lobule
−30	−11	54	55	Left precentral gyrus
−52	−38	21	54	Left inferior parietal lobule
50	−47	42	47	Right inferior parietal lobule
53	−47	24	43	Right supramarginal gyrus

BOLD amplitudes were then transformed to percent signal changes, relative to the mean BOLD amplitude calculated within the region across the run.

Note that because the MVPA classifier, although optimized, was associated with some degree of error in classification, resulting in a smoothing of the distribution of the demarcated shift points, we also computed

event-related average BOLD time courses for cued shifts of attention using the same algorithm. This was achieved by time-locking the BOLD signal to demarcated cued-shift points based on the output from the classifier when it was applied to the data from the cued-attention condition (using a leave-one-run-out procedure), rather than by time-locking to the actual onsets of the shift cues. Importantly, this method incorporates classifier error for the demarcation of both types of shifts (and avoids the need to correct for hemodynamic response lag at this stage), therefore allowing for a more appropriate and direct comparison between the uncued-shift and cued-shift event-related averages.

Finally, we performed a single planned contrast (i.e., the interaction between shift condition and time, one-tailed, α = .05, N = 12) on the event-related averages for each of these ROIs. Post-hoc tests of simple main effects were conducted following evidence for a reliable shift condition × time interaction; the statistical threshold for these post-hoc tests was Bonferroni-corrected (α = .025).
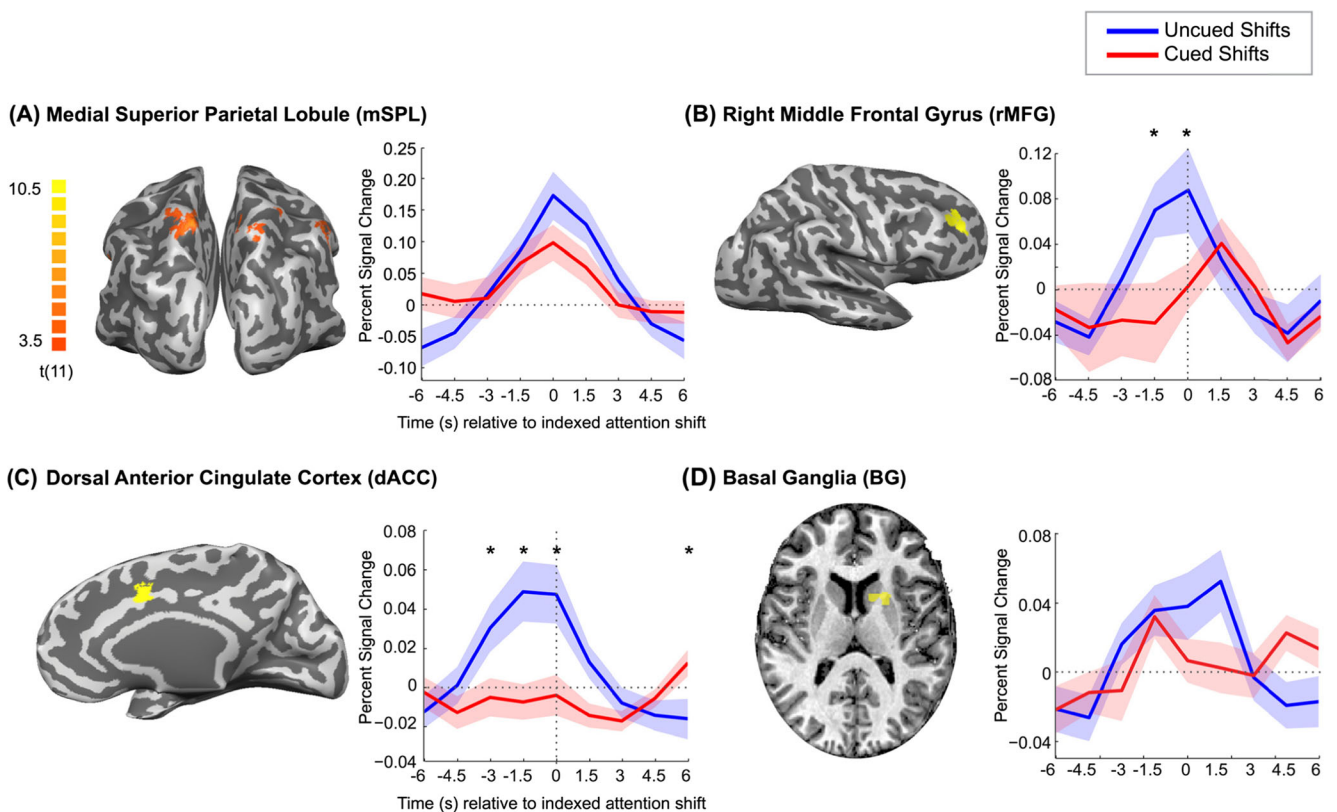

Fig. 2 Results. (A) Posterior view of regions reliably activated for cued attention shifts (based on a whole-brain Cued-Shift > Cued-Hold contrast; see Table 2), including mSPL, and the activation time courses within mSPL time-locked to attention shifts indexed by a multivoxel pattern analysis of visual cortex activity. (B–D) A priori regions of interest—rMFG, dACC, and BG (including left caudate nucleus and putamen; see Table 1)—and activation time courses. Shading indicates the standard errors of the means. For rMFG and dACC, asterisks indicate time points with reliable activation differences between the uncued (i.e., self-generated) and cued shifts (paired-samples t tests, p < .05). In the BG, we observed reliable uncued-shift activation, whereas the cued-shift activation and the interaction of time and type of shift (uncued vs. cued) did not reach significance

Results

Behavioral performance

In the *cued* condition of the fMRI session, response accuracy was 75.6 % correct on average (*SE* = 4.1 %). Note that chance-level accuracy of responses would have been only 25 % correct. Most of the errors that participants made were misses (i.e., failures to respond to the presence of a target; *M* = 18.6 %, *SE* = 3.3 %) rather than incorrect buttonpresses (*M* = 5.8 %, *SE* = 1.0 %). These results suggest that participants rarely, if ever, were attending to the wrong RSVP stream when responding to a target (otherwise, the rate of incorrect buttonpresses would have been comparatively high), consistent with the fact that the shift and hold cues (see the Method section) were uniquely colored (bright red), whereas the targets were not salient and required focused attention. In the *uncued* condition of the fMRI session, participants similarly missed 16.1 % of targets on average (*SE* = 3.4 %); there was no reliable difference in miss rates between the two conditions, $F(1, 11) = 1.57$, $p > .1$. Note that because in the uncued condition participants were free to attend to either RSVP stream and two different targets were presented simultaneously, response accuracy could not be calculated to enable a direct comparison across the task conditions. However, the classifier trained on data from the cued condition enabled us to predict participants' behavioral responses in the uncued condition with much greater than chance accuracy. Specifically, because attention could be classified as being directed toward the left or the right target at the time of target onset, this classification, together with the identity of the corresponding target, led to a prediction of the specific buttonpress that the participant would make. For example, if attention was classified as being directed toward the left target, which was the digit 4, then it was predicted that the participant would press the button corresponding to 4, rather than a button corresponding to the other possible targets: 2, 3, or 5. Using this simple binary classification, we correctly predicted participants' responses to 62.1 % of the targets, on average (whereas chance accuracy would be only 25 % correct).

Frequency of attention shifts in the cued condition

In the cued condition, the mean onset asynchrony between shift cues was 19.8 s (*SD* = 1.5 s; see the Method section). We first validated the classifier's ability to demarcate attention shifts within the cued condition. As expected, the MVPA model produced a mean onset asynchrony between the indexed cued shifts of attention of 18.49 s (*SD* = 0.8 s), which is within one TR of the actual shift-cue onset asynchrony. This result confirms that the classifier trained to discriminate Attend Left from Attend Right epochs during the cued condition was highly accurate in demarcating the transitions in

attention between left and right locations (i.e., attention shifts) in the cued condition for the test (left-out) runs. When the MVPA model was then applied to the data from the uncued condition, it produced a similar mean onset asynchrony between demarcated shifts of attention of 18.45 s (*SD* = 0.9 s), which is consistent with the instructions given to participants (i.e., to shift attention roughly three or four times per minute).

Neural mechanisms for self-generated, voluntary shifts of attention

We first replicated the primary finding of previous studies (Kelley et al., 2008; Yantis et al., 2002) that the mSPL was transiently active during cue-driven shifts of attention: A whole-brain Cued-Shift > Cued-Hold contrast (Table 2) yielded a suprathreshold cluster that was centered in left mSPL but extending into the right mSPL (Fig. 2A). We next extracted the event-related average time courses from this cluster, time-locked to the uncued attention shifts that were demarcated by the MVPA model. Transient activation in mSPL was revealed for uncued shifts [main effect of time: $F(8, 88) = 8.05$, $p < .001$], indicating that mSPL was engaged for both self-generated and cue-driven shifts of attention (Fig. 2A). An additional cluster in the left inferior parietal lobule (Table 2) that survived the whole-brain Cued-Shift > Cued-Hold contrast also clearly demonstrated transient activation time-locked to the demarcated uncued attention shifts [$F(8, 88) = 3.13$, $p = .004$], indicating that this region, like mSPL, was engaged for both self-generated and cue-driven shifts of attention. None of the remaining clusters in Table 2 exhibited reliable increases in activity time-locked to the uncued shifts of attention.

Within the rMFG (Fig. 2B), we found a reliable interaction between time and type of attention shift [$F(8, 88) = 1.79$, $p = .04$], potentially consistent with the hypothesis that self-generated shifts are associated with earlier rises in activity than are cue-driven shifts. Post-hoc tests of the simple main effects indicated that rMFG was transiently activated for uncued shifts [$F(8, 88) = 3.48$, $p = .002$], and although there appeared to be a trend for transient rMFG activity time-locked to the cued shifts, this effect was not reliable [$F(8, 88) = 1.05$, $p = .40$]. Similarly, within the dACC (Fig. 2C), a reliable interaction emerged between time and type of attention shift [$F(8, 88) = 5.70$, $p < .001$], potentially consistent with the hypothesis that self-generated shifts are associated with earlier rises in activity than are cue-driven shifts. Post-hoc tests of the simple main effects indicated that dACC was transiently activated for uncued shifts [$F(8, 88) = 6.87$, $p < .001$], but not for cued shifts [$F(8, 88) = 1.12$, $p = .36$].

Within the BG ROI (Fig. 2D), we observed transient activation time-locked to demarcated uncued shifts [effect of time on percent signal change: $F(8, 88) = 4.04$, $p < .001$]. The transient activation in the BG ROI time-locked to cued shifts

was not reliable [$F(8, 88) = 1.47, p = .18$]. The corresponding interaction approached statistical significance [$F(8, 88) = 1.83, p = .08$]. These results warrant the conclusion that the BG (left caudate nucleus and putamen) are transiently engaged for self-generated shifts of attention, but whether they are differentially engaged for cue-driven shifts of attention remains to be determined.

Discussion

Our first result of note is that mSPL is engaged not only for cue-driven shifts of attention, but also for self-generated shifts of attention. This finding was revealed by employing a novel method to identify self-generated, voluntary attention shifts occurring in the absence of instructional cues or overt responses that might otherwise have had confounding influences. Briefly, we first trained a classifier to discriminate multivoxel patterns of brain activity recorded following the presentation of instructional cues to attend to the left or the right location, over sustained epochs. We then applied this trained classifier to independent sets of data acquired while the participants freely shifted their attention between the two locations, to classify on a moment-by-moment basis each participant's focus of attention as being oriented to the left or the right location. Finally, we demarcated the time points at which the classifier indicated a leftward or a rightward shift of attention. During these uncued, self-generated shifts of attention, mSPL was transiently active (Fig. 2A), consistent with the hypothesis that it plays a role in reconfiguring attention. This finding also disconfirms the hypothesis that mSPL activation during cued shifts of attention merely reflects the processing of an external instruction to shift attention or the need to make an overt response.

In addition, we found that rMFG activity increased prior to self-generated attention shifts, and it increased earlier than the activity associated with cue-driven attention shifts (Fig. 2B). This finding suggests that rMFG participates in the preparation to reorient attention, much like premotor cortex participates in the preparation of overt movements (Wise, 1985). We also observed an early increase in dACC activity that was associated with self-generated, but not cue-driven, attention shifts (Fig. 2C). This result echoes evidence that medial frontal cortex participates in initiating or preparing self-generated, rather than externally triggered, overt actions such as hand movements (Krieghoff, Brass, Prinz, & Waszak, 2009; Passingham, Bengtsson, & Lau, 2010; Soon, Brass, Heinze, & Haynes, 2008). Together, these findings suggest that medial frontal cortex activity reflects a common locus of early-stage processing that gives rise to self-generated behavior, including both overt action and covert orienting of attention. We suggest that the dACC and rMFG are core components of the brain network underlying the will to act.

One might note that both of the a priori ROIs that demonstrated greater activity for self-generated than for cued shifts of attention prior to the demarcated attention shifts (rMFG and dACC) were regions that had been identified in an earlier study (Chiu et al., 2012) as being involved in cued shifts of attention. The evidence that these regions are involved in cued shifts of attention, however, was revealed in the earlier study only by MVPTC functional connectivity analyses, not by standard voxel-wise magnitude analyses. The present study indicates that, prior to the demarcated attention shifts, there is an increase in the magnitude of activation of these regions for self-generated shifts that is not observed reliably for cued shifts of attention, consistent with the earlier study. Chiu et al. also uncovered, via an MVPTC functional connectivity analysis, reliable activity in the BG related to cued shifts of attention that a standard method was unable to detect. The BG constitute a subcortical component within a network of functionally connected brain regions engaged for cue-driven shifts of attention (Chiu et al., 2012; Gitelman et al., 1999; Grande et al., 2006; Perry & Zeki, 2000; Shulman et al., 2009). In the present study, the BG (including clusters in the left caudate nucleus and putamen) exhibited transient activity related to self-generated attention shifts (Fig. 2D), a new finding that suggests a role of the BG in covert acts of volition. Together, the results of this and previous studies indicate that magnitude of activation, functional connectivity, and multivoxel patterns of activity may each reveal different roles of brain regions in both cued and self-generated shifts of attention, and they suggest that future studies should provide direct comparisons between these types of data to tease apart the roles of brain regions in different acts of attention.

In summary, our study introduces a novel approach to the investigation of willed behavior and cognition that monitors participants' ongoing brain activity, in the absence of instructional cues and overt responses, to index covert acts of volition—in this case, self-generated, voluntary shifts of visuospatial attention. This method revealed that mSPL, which was previously implicated in cue-driven shifts of attention, is also engaged in self-generated shifts of attention, a finding consistent with its hypothesized role in the reconfiguration of competitive interactions within sensory cortex (Yantis et al., 2002). Of particular note, regions in medial frontal and lateral prefrontal cortex (dACC, rMFG) exhibited early and selective increases in activation prior to self-generated shifts of attention. Activity within these regions likely reflects processing related to the intention or preparation to reorient attention. These results extend findings from investigations of other domains of cognitive control (e.g., Kühn et al., 2009; Schel et al., 2014). For example, a recent study (Schel et al., 2014) showed that whereas some brain regions were involved in both self-generated and cue-driven forms of response inhibition, self-generated inhibition additionally involved medial prefrontal cortex (Schel et al., 2014). The

present results also extend the findings of other studies (Soon et al., 2008) that have investigated freely chosen overt actions such as hand movements. Thus, our study provides critical new evidence for the core neural mechanisms underlying willed behavior in nonmotoric domains. It furthermore demonstrates the feasibility of an analysis method that can be adapted to investigate other domains of cognition (e.g., perceptual decision making and social judgments) that otherwise may be difficult or impossible to index in the absence of instructional cues or overt responses.

References

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10,* 433–436. doi:10.1163/156856897X00357

Chang, C. C., & Lin, C. J. (2011). LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology, 2*(3), 27.

Chiu, Y.-C., Esterman, M. S., Gmeindl, L., & Yantis, S. (2012). Tracking cognitive fluctuations with multivoxel pattern time course (MVPTC) analysis. *Neuropsychologia, 50,* 479–486. doi:10.1016/j.neuropsychologia.2011.07.007

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.

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience, 3,* 201–215.

Gitelman, D. R., Nobre, A. C., Parrish, T. B., LaBar, K. S., Kim, Y. H., Meyer, J. R., & Mesulam, M. (1999). A large-scale distributed network for covert spatial attention: Further anatomical delineation based on stringent behavioural and cognitive controls. *Brain, 122,* 1093–1106.

Gmeindl, L., Gao, J., Yantis, S., & Courtney, S. (2008, November). *Neural mechanisms of purely voluntary covert shifts of attention.* Paper presented at the Annual Meeting of the Society for Neuroscience, Washington, DC.

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. doi:10.1523/JNEUROSCI.4248-09.2010

Hopfinger, J. B., Camblin, C. C., & Parks, E. L. (2010). Isolating the internal in endogenous attention. *Psychophysiology, 47,* 739–747.

Kelley, T. A., Serences, J. T., Giesbrecht, B., & Yantis, S. (2008). Cortical mechanisms for shifting and holding visuospatial attention. *Cerebral Cortex, 18,* 114–125.

Krieghoff, V., Brass, M., Prinz, W., & Waszak, F. (2009). Dissociating what and when of intentional actions. *Frontiers in Human Neuroscience, 3,* 3. doi:10.3389/neuro.09.003.2009

Kühn, S., Haggard, P., & Brass, M. (2009). Intentional inhibition: How the "veto-area" exerts control. *Human Brain Mapping, 30,* 2834–2843. doi:10.1002/hbm.20711

Lau, H. C., Rogers, R. D., & Passingham, R. E. (2006). On measuring the perceived onsets of spontaneous actions. *Journal of Neuroscience, 26,* 7265–7271.

Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential): The unconscious initiation of a freely voluntary act. *Brain, 106,* 623–642.

Passingham, R. E., Bengtsson, S. L., & Lau, H. C. (2010). Medial frontal cortex: From self-generated action to reflection on one's own performance. *Trends in Cognitive Sciences, 14,* 16–21. doi:10.1016/j.tics.2009.11.001

Perry, R. J., & Zeki, S. (2000). The neurology of saccades and covert shifts in spatial attention: An event-related fMRI study. *Brain, 123,* 2273–2288.

Schel, M. A., Kühn, S., Brass, M., Haggard, P., Ridderinkhof, K. R., & Crone, E. A. (2014). Neural correlates of intentional and stimulus-driven inhibition: A comparison. *Frontiers in Human Neuroscience, 8*(27), 1–10. doi:10.3389/fnhum.2014.00027

Serences, J. T. (2004). A comparison of methods for characterizing the event-related BOLD timeseries in rapid fMRI. *NeuroImage, 21,* 1690–1700.

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. doi:10.1093/cercor/bhh095

Shomstein, S., & Yantis, S. (2004). Control of attention shifts between vision and audition in human cortex. *Journal of Neuroscience, 24,* 10702–10706.

Shulman, G. L., Astafiev, S. V., Franke, D., Pope, D. L. W., Snyder, A. Z., McAvoy, M. P., & Corbetta, M. (2009). Interaction of stimulus-driven reorienting and expectation in ventral and dorsal frontoparietal and basal ganglia-cortical networks. *Journal of Neuroscience, 29,* 4392–4407.

Slotnick, S. D., Serences, J. T., & Yantis, S. (2003). Cortical mechanisms of feature-based attentional control. *Cerebral Cortex, 13,* 1334–1343.

Soon, C. S., Brass, M., Heinze, H.-J., & Haynes, J.-D. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience, 11,* 543–545.

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, 23,* 2905–2919. doi:10.1162/jocn.2011.21608

Taylor, P. C. J., Rushworth, M. F. S., & Nobre, A. C. (2008). Choosing where to attend and the medial frontal cortex: An FMRI study. *Journal of Neurophysiology, 100,* 1397–1406.

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.

Wise, S. P. (1985). The primate premotor cortex: Past, present, and preparatory. *Annual Review of Neuroscience, 8,* 1–19.

Yantis, S., Schwarzbach, J., Serences, J. T., Carlson, R. L., Steinmetz, M. A., Pekar, J. J., & Courtney, S. M. (2002). Transient neural activity in human parietal cortex during spatial attention shifts. *Nature Neuroscience, 5,* 995–1002. doi:10.1038/nn921