

Decoding the Content of Delayed Intentions

Sam J. Gilbert

Institute of Cognitive Neuroscience and Division of Psychology and Language Sciences, University College London, London WC1N 3AR, United Kingdom

Rostrolateral prefrontal cortex (RLPFC) plays a key role in our ability to postpone the execution of intended behaviors until after another activity has been performed. However, it is poorly understood in computational terms. One crucial question is whether RLPFC represents the content of delayed intentions or plays a nonspecific role. In this human functional magnetic resonance imaging study ($n = 32$), RLPFC was active while participants stored delayed intentions during a distracting ongoing task. Multivariate analysis showed that the intended cue for future action and the intended behavior could be decoded from distinct posterior brain regions. However, the content of intentions could not be decoded from RLPFC itself. Functional connectivity analysis showed that RLPFC increased its coupling with content-representing regions during intention storage. Furthermore, trials with relatively high RLPFC activity were associated with enhanced decoding. Thus, RLPFC may enable realization of delayed intentions via interactions with posterior brain regions, which represent their content.

Introduction

Rostrolateral prefrontal cortex (RLPFC), corresponding to the lateral aspects of Brodmann area (BA) 10, is widely considered to occupy the top of a hierarchy of lateral prefrontal regions supporting behavioral regulation (Badre and D'Esposito, 2009; Christoff et al., 2009; Koechlin et al., 2003). Neuroimaging and neuropsychological evidence implicates RLPFC in our ability to hold an overarching goal in mind while engaged in subordinate tasks: "cognitive branching" (Koechlin et al., 1999) or "goal-subgoal integration" (Braver and Bongiolatti, 2002). This allows us to pursue complex, hierarchically organized goals, for example when we need to hold in mind a delayed intention, to be realized at a later time. This is sometimes termed "prospective memory" (PM) (Brandimonte et al., 1996; McDaniel and Einstein, 2000; Kliegel et al., 2008; West, 2008). Accordingly, patients with RLPFC damage can exhibit a selective impairment in PM (Burgess et al., 2000; Uretzky and Gilboa, 2010).

Despite clear evidence denoting a class of situations that recruit RLPFC, the precise role of RLPFC is poorly understood in information processing terms. One crucial question is whether RLPFC itself represents the content of delayed intentions, or alternatively plays a non-content-specific role that supports PM without representing specific intentions. According to the first of these hypotheses, RLPFC can be thought of as implementing a "goal stack," representing the content of goals to be achieved. Alternatively, RLPFC might be thought to support a state promoting the ability of environmental events to act as retrieval cues for intended behaviors, analogous to "retrieval mode" (Tulving, 1983). According to this latter model, RLPFC might play a non-

content-specific role via its interactions with distinct brain regions that themselves represent the content of intentions (Sakai and Passingham, 2003; Rowe et al., 2007).

To investigate this question, the present study used multivoxel pattern analysis (MVPA) to decode participants' intentions from patterns of brain activity. Insofar as a delayed intention may be decoded from a particular brain region, this indicates a representation of that intention in the relevant brain area. Furthermore, this study examined whether different aspects of participants' intentions could be decoded from distinct brain regions. According to some authors (Gollwitzer, 1999), intentions can often be conceptualized as mental links between an anticipated future cue (e.g., seeing a particular person at some future time) and a specific behavior (e.g., passing on a particular message). Such intentions require participants to represent (1) the future cue to detect and (2) the appropriate response that should then be produced. These two aspects were examined separately, to investigate the overlap (or lack of it) between brain regions involved in their representation.

An additional unresolved issue concerns the processing stage at which RLPFC is involved in PM: (1) encoding intentions into memory, (2) maintaining those intentions during ongoing-task performance, or (3) retrieving those intentions at the appropriate time. Although previous studies suggest that simply maintaining intentions is sufficient to yield RLPFC activation (Burgess et al., 2001; Reynolds et al., 2009), the extent to which RLPFC also plays a role in encoding and retrieving intentions is unclear.

Materials and Methods

Participants

Thirty-two right-handed participants took part in the study (18 males; mean age, 25; range, 19–38). The experiment was approved by the local ethics committee, and all participants provided written informed consent before taking part.

Behavioral task

A paradigm was developed in which participants repeatedly (1) formed, (2) maintained, and then (3) realized delayed intentions, while engaged in an unrelated ongoing task. By pseudorandomly manipulating the du-

Received Oct. 11, 2010; revised Dec. 16, 2010; accepted Dec. 21, 2010.

This study was supported by a Royal Society University Research Fellowship. I am extremely grateful to Diana Armbruster and Maria Panagiotidi for assistance with data collection, to staff at the Birkbeck–University College London Centre for Neuroimaging for assistance with scanning, and to Roland Benoit and Paul Burgess for helpful comments on a previous draft of this article.

Correspondence should be addressed to Sam J. Gilbert, Institute of Cognitive Neuroscience, 17 Queen Square, London WC1N 3AR, UK. E-mail: sam.gilbert@ucl.ac.uk.

DOI:10.1523/JNEUROSCI.5336-10.2011

Copyright © 2011 the authors 0270-6474/11/312888-07\$15.00/0

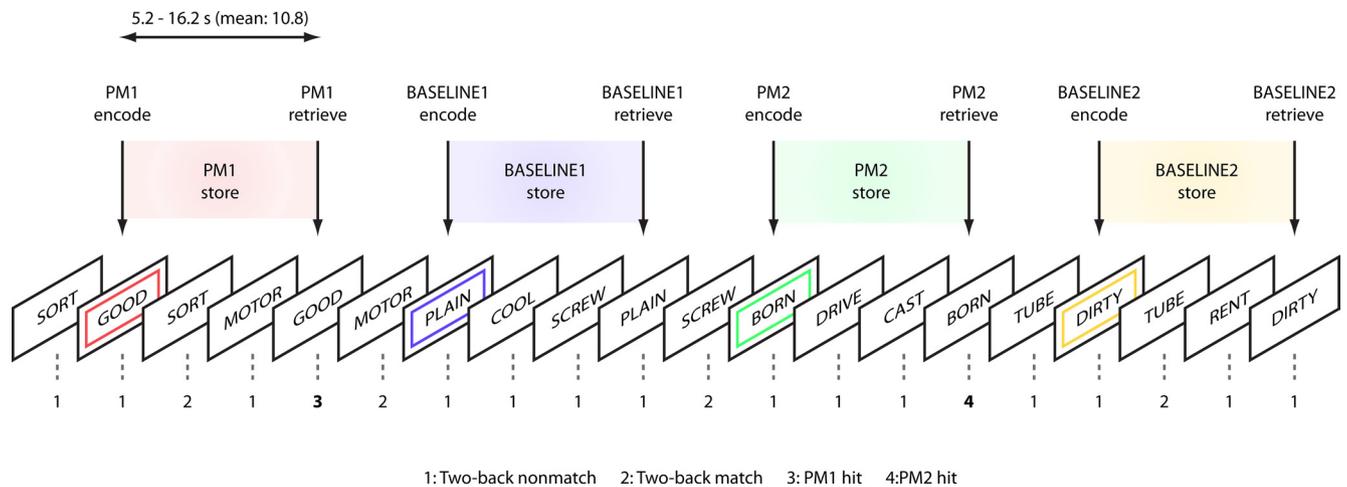


Figure 1. Schematic illustration of the task. Participants performed a two-back task (responding with button 1 to two-back nonmatches and button 2 to two-back matches). Encode–store–retrieve cycles were embedded within this task. Stimuli were surrounded by colored borders on encode trials, with the assignment of colors to conditions counterbalanced between participants [e.g., red: PM1 encode (press button 3 when you encounter the stimulus again); green: PM2 encode (press button 4 when you encounter the stimulus again); blue and yellow: baseline 1 and 2 (no additional task requirements)]. PM and baseline cycles alternated. Version 1 of the task included PM1 and baseline 1 conditions only. In version 2 of the task, the selection of PM1 versus PM2 for a particular PM cycle, and baseline1 versus baseline2 for a particular baseline cycle, was pseudorandomized. In half of the blocks, stimuli consisted of words, as illustrated; in the other half of blocks, picture stimuli were used.

ration of the maintenance period, it was possible to disambiguate these three phases (cf. Rowe et al., 2000) (Fig. 1). Participants performed a two-back task throughout the experiment, indicating on each trial whether the stimulus matched the item presented two trials previously (interstimulus interval, 2.7 s). A PM task was embedded within this two-back task (Reynolds et al., 2009). The stimuli consisted of pictures in half of the blocks, and words in the other half. On “PM encode trials,” the stimulus was surrounded by a colored border (red or blue, counterbalanced across participants). On such trials, participants performed the two-back task as usual. However, they were asked to remember the stimulus, and to press a third button, rather than one of the buttons used for the two-back task, should they notice the same stimulus on a subsequent trial. Between two and six trials after each PM encode trial, the relevant stimulus was repeated (“PM retrieve trials”), after which it was not shown again. Thus, on PM retrieve trials, participants were required to interrupt their ongoing performance of the two-back task and instead press a different button to indicate that they had noticed the stimulus that they had encoded on the previous PM encoding trial. Crucially, the stimuli presented on PM retrieve trials were no different in appearance to those presented on other trials, and a two-back response to such stimuli would be possible if they were not detected as targets. Correct responses on PM retrieve trials therefore had to be self-initiated on the occurrence of the appropriate item, rather than being strongly cued by the stimulus characteristics. This is one of the defining features of laboratory paradigms assessing PM, as opposed to other forms of dual-task paradigm (Burgess et al., 2003). The trials between PM encoding and PM retrieval trials were denoted “PM store trials.” During these trials, participants had to maintain a representation of the stimulus presented on the PM encode trial so that they could interrupt their performance of the two-back task when it was presented again on a PM retrieve trial.

On “baseline encode trials,” the stimulus was also surrounded by a different colored border (blue/red counterbalanced across participants), but participants were instructed to treat these trials identically to any other without a border. These baseline encode trials were followed by “baseline store trials” and “baseline retrieve trials” in an identical manner to the PM trials, but there were no task requirements associated with these trials apart from the two-back task. Thus, baseline trials were matched perceptually with PM trials, but they did not require encoding, storage, or retrieval of delayed intentions.

Two different versions of the tasks were performed by two groups of 16 participants. The first version of the task was as above. In the second version of the task, there were four possible colored borders (red, green, blue, yellow) and two PM response buttons instead of one. In this version

of the task, two colors instructed participants to press one of two PM response buttons on the subsequent PM retrieval trial; the other two colors served as baselines. Thus, in this version, participants’ delayed intentions varied both in the nature of the cue to detect (word or picture) and the response to make (third or fourth finger). The first, simplified, version of the task was administered so as to keep the number of cue to response mappings to a minimum for multivariate decoding. However, decoding accuracy did not differ significantly between the two versions of the task. Results from the two task versions were therefore collapsed.

Behavioral procedure

Participants performed the task for six sessions, three featuring word stimuli and the other three featuring picture stimuli. Material type (words or pictures) alternated with each session and was counterbalanced across participants for the first session. In each session, participants viewed a series of 123 stimuli at the center of a screen, viewable via a mirror. Pictures were randomly selected images from the photographic collection “20,000 Photos” (Focus Multimedia Limited) including animals, cityscapes, landscapes, buildings, and vehicles, at a size of ~5° tall and wide. Words were randomly selected from the MRC Psycholinguistic database with the constraint that they were three to five letters long, associated with a familiarity rating of 400–600 and an imagability rating of 400–700. Words were presented in capital letters, Arial typeface, ~1° tall. On each trial, the stimulus was presented for 1.35 s, followed by a blank screen for 1.35 s. Participants could respond at any point in the trial.

To select stimuli for each session, random items were first assigned to each of the 123 trials. After this, on a randomly selected 50% of trials, the item was replaced with the stimulus that was presented two trials previously (i.e., these trials were made into two-back hits). The PM and baseline encoding and retrieval trials were then inserted. The first encoding event occurred on the second trial of each session, with the nature of the trial (PM or baseline) randomly chosen. On these trials, the stimulus was surrounded by a colored border (red or blue in version 1 of the task; red, blue, green, or yellow in version 2 of the task; the assignment of colors to conditions was counterbalanced across participants). After this encoding event, there were between one and five store trials, selected from a uniform distribution. The same stimulus that was presented in the encoding trial was then presented, followed by one filler trial and then the next encoding trial. Successive encoding trials always alternated between PM and baseline conditions. In version 2 of the task, the two categories of PM and baseline conditions were randomly selected, with the constraint that they occurred in equal number within every session and no more than three in a row were of the same category (e.g., PM1, baseline2, PM1,

baseline1, PM1, baseline2 would necessarily be followed by a PM2 event). A total of 20 encode–store–retrieve cycles were presented in each session, preceded by one filler trial (the first trial of the session) and followed by two filler trials (the last two trials of the session) to yield a total of 123 trials in each session: three filler trials plus 20 cycles consisting of an encoding trial, between one and five store trials (mean, three), a retrieval trial, and one filler trial (i.e., 20 cycles averaging six trials each).

MRI scanning

A 1.5 T Siemens TIM Avanto scanner was used to acquire both T1-weighted structural images and T2*-weighted echoplanar images (64×64 ; 3.5×3.5 mm pixels; echo time, 40 ms) with blood oxygen level-dependent (BOLD) contrast. Each volume comprised 31 axial slices (3.5 mm thick, oriented approximately to the anterior commissure–posterior commissure plane). Functional scans were acquired in six sessions, each comprising 135 volumes (~6 min). Volumes were acquired continuously with an effective repetition time of 2.5 s per volume. The first four volumes in each session were discarded to allow for T1 equilibration effects. Between the third and fourth functional scan, a 6 min T1-weighted structural scan was performed.

Data analysis

Univariate analyses. Functional magnetic resonance imaging data were analyzed using SPM8 software (<http://www.fil.ion.ucl.ac.uk/spm/software/spm8/>). The volumes were realigned, corrected for different slice acquisition times, normalized into 3 mm cubic voxels using the Montreal Neurological Institute reference brain and fourth-degree B-spline interpolation, and smoothed with an isotropic 8 mm full-width half-maximum Gaussian kernel. The volumes acquired during the six sessions were treated as separate time series. For each series, the variance in the BOLD signal was decomposed with a set of regressors in a general linear model (Friston et al., 1995), using a mixed blocked and event-related design (Visscher et al., 2003). Separate regressors coded for (1) PM encode, (2) PM retrieve, (3) baseline encode, (4) baseline retrieve events, using delta functions aligned to the onset of each event. In version 2 of the experiment (with two different colors associated with the PM and baseline conditions), PM and baseline events were each coded with two separate regressors. Additional boxcar regressors coded for (1) PM storage and (2) baseline storage periods (again with additional regressors in version 2 of the experiment), indicating the period between relevant encoding and retrieval events (mean duration, 10.8 s; range, 5.2–16.2 s). By pseudorandomly manipulating the duration of the storage periods, activity associated with these periods was distinguished from activity associated with the encode and retrieve events (Rowe et al., 2000; Visscher et al., 2003). All regressors were convolved with a canonical hemodynamic response function. These regressors, together with the regressors representing residual movement-related artifacts and the mean over scans, comprised the full model for each session. The data and model were high-pass filtered to a cutoff of 1/128 Hz.

Parameter estimates for each regressor were calculated from least mean squares fit of the model to the data. Effects of interest were assessed in a random-effects analysis as follows. Separate analyses were conducted investigating each of the three phases (encoding, storage, retrieval). For each phase, contrast images representing conditions of interest were entered into a repeated-measures ANOVA using nonsphericity correction (Friston et al., 2002). The four conditions of interest entered into this analysis represented the factorial crossing of material type (words, pictures) and condition (PM, baseline) (e.g., for the analysis of encoding events, the conditions of interest were PM encode word, PM encode picture, baseline encode word, baseline encode picture). Appropriate contrasts for effects of interest were conducted at the second level, applying a height threshold of $p < 0.001$ uncorrected in conjunction with an extent threshold determined by SPM8 to achieve $p < 0.05$ familywise error correction for multiple comparisons across the whole brain volume.

MVPA. A searchlight approach was used (Kriegeskorte et al., 2006), investigating decoding accuracy from a sphere of voxels centered on each voxel in the brain in turn (radius, three voxels). Two orthogonal aspects of PM intentions were decoded, in separate analyses: (1) whether the forthcoming cue for a PM response was a word or picture; (2) which of

the two PM responses was required: third or fourth finger (for participants performing version 2 of the task). These analyses were restricted to the storage period between the encoding and retrieval trials. Seeing as word and picture stimuli were presented in separate blocks, it could be trivial to decode whether the forthcoming PM cue is a word or a picture, seeing as this might simply reflect whether the current ongoing stimuli are words or pictures. To exclude this possibility, all MVPA analyses searched for significant differences in decoding accuracy between PM storage periods and baseline storage periods. Since the PM and baseline periods alternated within each block, this allowed the detection of PM-specific effects.

In these analyses, preprocessing did not include normalization or smoothing. First-level models were generated for each participant in SPM8 as above. However, instead of two regressors coding for PM store and baseline store periods (or four regressors in version 2 of the task), every store period was modeled with a separate regressor [i.e., a total of 120 regressors (20 store periods per session, divided equally between PM and baseline conditions)]. This yielded 120 images per participant, consisting of parameter estimates at each voxel corresponding to each of the 120 store periods (60 PM and 60 baseline). Half of these 120 were derived from sessions using word stimuli and the other half from sessions using picture stimuli. In version 2 of the experiment, orthogonal to the word/picture distinction, the PM and baseline conditions were subdivided into two versions associated with different colored borders in the preceding encoding trial. MVPA results were calculated using these 120 images as raw data. At each voxel, a spherical region of interest (ROI) was generated (radius, three voxels). Parameter estimates for each voxel within the ROI were extracted, for each of the 120 storage periods. This yielded 120 vectors, each representing a set of parameter estimates at a particular ROI for a specific store period. Each vector was normalized to mean 0, SD 1, so that decoding accuracy was based on the distribution of activation over voxels rather than differences in mean level of activation.

First, the 60 PM store periods were examined. In the cue decoding analysis, an attempt was made to distinguish the 30 PM store periods using word stimuli versus the 30 PM store periods using picture stimuli. In the response decoding analysis (version 2 of the task only), an attempt was made to distinguish the 30 PM store periods requiring a third-button response on cue detection, versus the 30 PM store periods requiring a fourth-button response on cue detection. For each of the two categories of PM trial, the first exemplar was discarded and a support vector machine trained on the remaining 29 exemplars of each type (LIBSVM implementation, <http://www.csie.ntu.edu.tw/~cjlin/libsvm/>; regularization parameter, 1). The support vector machine was then tested on the remaining two exemplars, one of each type. This procedure was then repeated after rotating the training/testing sets over the remaining stimuli, and the mean classification accuracy was recorded. After this, the procedure was repeated using the baseline store periods (e.g., comparing baseline store periods using word vs picture stimuli). Accuracy for baseline store periods was subtracted from PM store periods. The resulting maps, indicating the difference between PM and baseline decoding accuracy at each voxel, were then normalized into Montreal Neurological Institute space (using the same normalization parameters as the univariate analyses) and smoothed using a Gaussian kernel (full-width half-maximum, 4 mm). This relatively small kernel size was used to avoid excessive smoothing, seeing as the searchlight analysis already imposes spatial smoothing on the data. The resulting images were entered into a one-sample *t* test using SPM8. This allowed assessment of voxels showing consistently higher decoding accuracy for PM than baseline store periods, across participants, in a random-effects analysis.

Along with the main MVPA analyses described above, two additional analyses were conducted. In the first analysis, trials with relatively high versus low RLPFC activation were compared. First, a mask was generated of all voxels activated in the group-level univariate contrast of PM store > baseline store at an uncorrected threshold of $p < 0.001$, restricted to the approximate region of Brodmann area 10 (using the mask supplied with WFU Pickatlas; <http://fmri.wfubmc.edu/cms/software/>). Store periods in each category were then divided into two halves, based on a median split according to signal intensity within this mask associated with the relevant store period. The MVPA analysis was then conducted as above,

separately for store periods associated with high versus low RLPFC signal intensity. Each participant's results from low RLPFC signal intensity trials were subtracted from results from high RLPFC signal intensity trials. The resulting images were entered into one-sample *t* tests as above, to test for regions showing significantly greater intention decoding accuracy on high-RLPFC signal trials than low-RLPFC signal trials.

In a final MVPA analysis, a classifier was trained to distinguish PM store from baseline store periods, separately for word and picture stimuli. In this analysis, training and testing data were always derived from separate sessions (i.e., the classifier was trained on data from two sessions and tested on data from the third, with the training and testing sets rotated over the three sessions for each stimulus type). This ensured that accuracy could not be inflated by non-independent training and testing data, derived from the same session. Note that this issue could not affect the other MVPA analyses described above, seeing as they reflect differences in accuracy between PM and baseline store periods, both of which were subjected to the same cross-validation technique (see supplemental information, available at www.jneurosci.org as supplemental material). Resulting accuracy maps from these MVPA analyses were normalized and smoothed as above. These images were then entered into a repeated-measures ANOVA (within-subject factor: stimulus type) as in the univariate analyses, for assessment of results at the group level.

Psychophysiological interaction analysis. The mask described above, indicating regions within RLPFC showing significant activation for the PM store > baseline store contrast, was used as the seed region for a psychophysiological interaction (PPI) analysis (Friston et al., 1997). Voxelwise PPI analysis was conducted to identify brain regions showing a significant increase in functional coupling with this seed region during PM store versus baseline store periods. A PPI model was created for each participant, using the BOLD time course extracted from the seed region to generate the physiological variable and the contrast of PM store versus baseline store to create the psychological variable. At the first level, the PPI model included the physiological variable, psychological variable, and PPI, along with movement regressors and the mean over scans. Contrast images associated with the PPI regressors were entered into a one-sample *t* test, using a height threshold of $p < 0.001$, uncorrected, cluster-corrected to achieve a whole-brain familywise error-corrected threshold of $p < 0.05$. Analyses of the PPI results at regions defined by the MVPA results were performed using small-volume corrections based on 10-mm-radius spheres centered on the peak MVPA coordinates, applying a familywise error correction across voxels included in the region of interest.

Results

Behavioral results

Behavioral results did not differ significantly between the two versions of the task on any measure and were collapsed. Accuracy in the two-back task was relatively high, marginally higher for words than pictures (words, 88%; pictures, 87%; $F_{(1,30)} = 3.3$, $p = 0.08$). Mean response times (RTs) to words were marginally faster than to pictures in this task (words, 850 ms; pictures, 870 ms; $F_{(1,30)} = 2.95$, $p = 0.096$). Participants accurately pressed the PM response button on 67% of PM retrieval trials, responding more accurately for picture (70%) than word (64%) trials ($F_{(1,30)} = 5.9$; $p = 0.02$). In version 2 of the task, response confusions (i.e., pressing the wrong PM response button on a PM retrieval trial) were rare (4% of PM retrieval trials). PM accuracy in the present study was therefore comparable with previous studies of PM (Simons et al., 2006; Gilbert et al., 2009) despite a mean delay between PM encoding and the opportunity to respond of just 10.8 s, much shorter than previous studies (typically on the order of several minutes or more). This suggests that the two-back task was sufficiently distracting to prevent continuous rehearsal of PM intentions (Einstein et al., 2003).

RTs on PM storage trials were significantly slower than baseline storage trials (PM, 878 ms; baseline, 842 ms; $F_{(1,30)} = 29.2$, $p < 0.0001$), indicating a RT cost associated with maintaining PM intentions and/or monitoring ongoing trials for PM cues (Smith,

2003; Einstein et al., 2005). This effect was stronger for word stimuli than picture stimuli (word PM, 877 ms; word baseline, 824 ms; picture PM, 879 ms; picture baseline, 860 ms; $F_{(1,30)} = 28.4$, $p < 0.001$). However, PM-related RT slowing was uncorrelated with MVPA results (see supplemental information, available at www.jneurosci.org as supplemental material). Finally, RTs during PM storage periods were separated according to whether they were followed by a correct PM response ("PM hit") or not ("PM miss"). RTs during storage periods associated with both hits and misses (873 and 884 ms, respectively) were significantly slower than RTs during baseline storage periods (842 ms; $F_{(1,30)} > 18.2$, $p < 0.001$). However, RTs did not differ significantly between PM storage periods associated with subsequent hits versus misses ($F_{(1,30)} = 2.0$; $p = 0.166$). This suggests that participants encoded and maintained intentions at least on a proportion of trials associated with a subsequent PM miss.

Univariate analyses of neuroimaging data

In the neuroimaging analysis, PM encoding trials were first compared against baseline encoding trials, to identify regions involved in encoding PM intentions. This revealed activity in widespread occipital, medial and lateral temporal, medial and lateral parietal, and subcortical regions (Fig. 2; supplemental information, available at www.jneurosci.org as supplemental material). Apart from a small region of left dorsolateral prefrontal cortex (BA 46), the absence of activation within prefrontal cortex for this contrast was striking. Given the well known role of the hippocampus in memory encoding (Eichenbaum and Cohen, 2001; Simons and Spiers, 2004; Squire et al., 2004), and suggestions that the medial temporal lobes play a role in representing PM intentions (Adda et al., 2008), it is notable that bilateral hippocampal activation was observed for this contrast (Fig. 2). PM storage trials were then compared against baseline storage trials, to identify regions involved in storing PM intentions and/or monitoring for PM cues. This revealed strong bilateral activation in lateral prefrontal cortex, centered on BA 10/46 (RLPFC), extending into bilateral insula and anterior cingulate. Additional activation was seen in parietal cortex and cerebellum. Finally, PM retrieval trials were compared against baseline retrieval trials, to identify regions involved in retrieving and realizing PM intentions. This revealed activity in bilateral precentral and postcentral gyrus, bilateral cerebellum, and posterior cingulate. These results are shown in Figure 2. All results remained similar (i.e., activations remained significant) when an alternative analysis was performed, limited to correct PM trials only.

Notably, only the PM storage period, and not the encoding and retrieval events, was associated with significant activation in RLPFC. To confirm the specificity of RLPFC activity for PM storage versus encoding and retrieval, signal was extracted from the peak left and right RLPFC regions identified by the PM > ongoing-only contrasts reported by Simons et al. (2006) (left, $-39\ 57\ 3$; right, $39\ 54\ 15$). In both left and right RLPFC, the contrast of PM storage versus baseline storage revealed significant activation ($t_{(31)} > 2.7$; $p < 0.011$). However, neither the encoding-related nor the retrieval-related contrasts yielded significant activation in either region ($t_{(31)} < 1.33$; $p > 0.19$).

MVPA

The present results confirmed the involvement of RLPFC in PM, specifically while participants maintained a future intention and monitored their environment for a PM cue, rather than when they encoded or realized their intentions. Next, MVPA was used in an attempt to decode the content of participants' intentions on

a trial-by-trial basis. The cue decoding analysis revealed four brain regions from which it was possible to decode whether participants were anticipating a word versus a picture cue for future action: medial rostral prefrontal cortex, thalamus, right superior parietal cortex, and medial occipital cortex (Fig. 3, Table 1). However, there was no evidence that RLPFC represented the content of participants' intentions. Even after performing small volume corrections on regions of interest defined by 10-mm-radius spheres centered on peak RLPFC coordinates from the univariate contrast of PM storage versus baseline storage, there was still no significant cue decoding effect associated with RLPFC.

The response decoding analysis revealed more widespread brain regions from which it was possible to decode which of the two PM responses had to be made, including motor cortex, putamen, medial temporal lobe, and lateral occipital cortex; these regions were remarkably distinct from those identified in the cue decoding analysis (Fig. 3, Table 1). However, again, there was no evidence that RLPFC represented the content of participants' intentions, even after small volume corrections. Analogous MVPA analyses of cue and response decoding for encode and retrieve trials, rather than store periods, also failed to reveal significant effects in RLPFC.

Might these results simply indicate a lack of power to detect an effect in RLPFC, using an MVPA approach? To answer this question, an attempt was made to decode whether each store period belonged to the PM or baseline condition. This revealed a highly significant effect in bilateral RLPFC (peak coordinates: $-33, 50, -2$, $Z_{\max} = 6.15$; $33, 53, 13$, $Z_{\max} = 5.95$). Thus, it was possible to decode from activity within RLPFC whether or not an intention was being held, but not the content of that intention. This shows that the absence of RLPFC effects in the cue and response decoding analyses do not merely indicate an inability to detect such effects using MVPA. Note that overall levels of activity were normalized for all MVPA analyses, so these results cannot reflect different mean activity levels for PM versus baseline, as revealed in the univariate analysis.

Functional connectivity analysis

These results present a conundrum. On the one hand, univariate analyses indicated strong engagement of RLPFC while participants stored PM intentions, consistent with previous neuroimaging and neuropsychological studies indicating the involvement of this region in PM (Okuda et al., 1998; Burgess et al., 2000, 2001, 2003; Simons et al., 2006; West, 2008; Gilbert et al., 2009; Reynolds et al., 2009; Uretzky and Gilboa, 2010). However, it was not possible to decode the content of those intentions from patterns of activity within RLPFC (unlike widespread posterior brain regions). The role

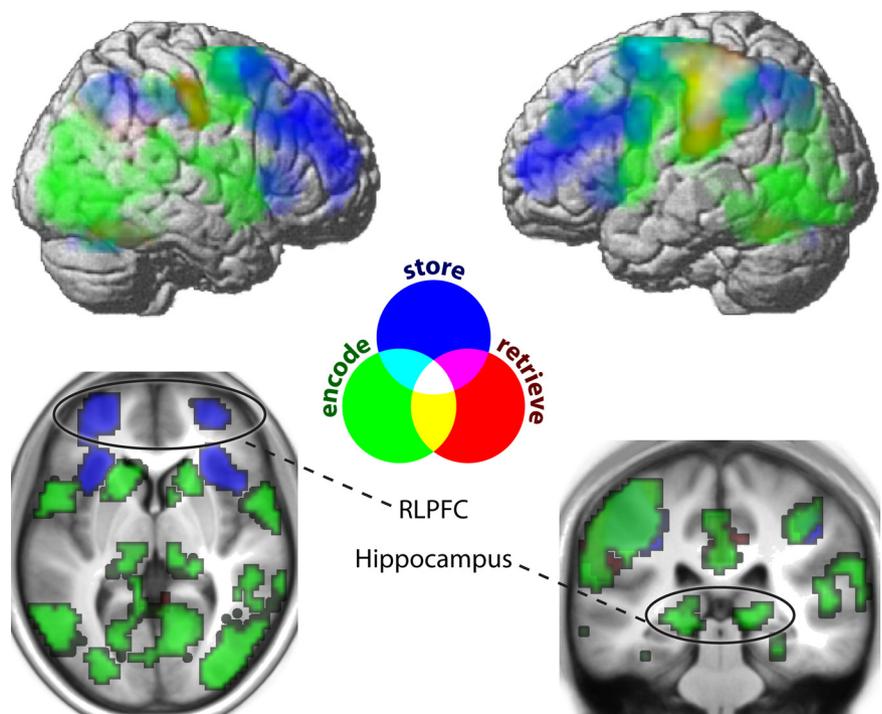


Figure 2. Regions of significant activation in the univariate comparisons of PM versus baseline conditions, presented separately for encode, store, and retrieve periods. Results are presented on a three-dimensional rendering of a template brain (top) and axial ($z = 4$) and coronal ($y = -32$) sections of the mean normalized structural scan (bottom).

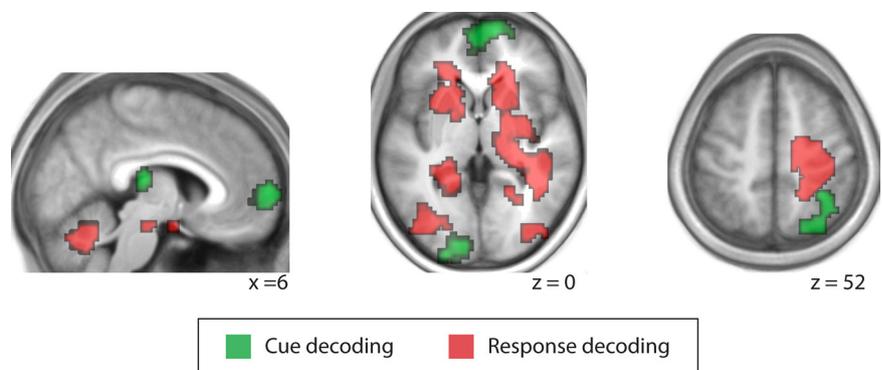


Figure 3. Regions from which it was possible to decode the forthcoming PM cue (word vs picture) or the forthcoming PM response (third vs fourth finger) during PM store periods with significantly greater accuracy than the corresponding baseline storage periods. Results are presented on sections of the mean normalized structural scan.

of RLPFC in PM is therefore unclear. As hypothesized above, one possibility is that RLPFC supports PM via its interactions with other brain regions, rather than itself representing the content of delayed intentions (Sakai and Passingham, 2003; Rowe et al., 2007). To test this hypothesis, a PPI analysis was conducted, to search for brain regions showing greater functional coupling with RLPFC during PM store periods than baseline store periods, defining a seed region as all those voxels within Brodmann area 10 that were activated by the contrast of PM store periods versus baseline periods at a threshold of $p < 0.001$, uncorrected, in the group-level analysis.

At a whole-brain corrected threshold, this PPI analysis revealed several posterior regions showing enhanced coupling with RLPFC during PM store periods (see supplemental information, available at www.jneurosci.org as supplemental material). Importantly, using small volume corrections based on 10-mm-radius spheres centered on the peak coordinates from the MVPA

Table 1. MVPA results

Region	BA	Coordinate	Decoding accuracy (%)	Z_{\max}	N voxels	PPI
Cue decoding						
Medial rostral PFC	10	9, 62, 4	3.4	4.35	170	0.055
Thalamus	—	6, -22, 13	2.9	4.30	75	0.012
Medial occipital cortex	17	-18, -88, -5	1.8	4.21	100	NS
Lateral parietal cortex	7/40	36, -49, 52	2.7	4.08	137	NS
Response decoding						
Medial parietal cortex	7/31	21, -37, 34	9.1	4.93	23,730	NS
Lateral occipital cortex	19/37	-39, -76, 1	9.0	4.62	222	0.074
Putamen	—	-27, 14, -2	7.4	4.37	857	0.023
Lateral occipital cortex	19	36, -82, 1	8.2	4.01	68	0.025
Medial parietal cortex	23	-18, -10, 43	6.7	3.61	61	NS

The table shows peak coordinates of regions from which it was possible to decode the forthcoming PM cue (word vs picture) or the forthcoming PM response (third vs fourth finger) during PM storage periods with significantly greater accuracy than the corresponding baseline storage periods. Decoding accuracy is baseline corrected, so that 0% indicates equivalent decoding accuracy between PM and baseline storage periods. Thus, although accuracy appears greater for response than cue decoding, these two types of figure are not directly comparable, seeing as baselines differed between word and picture storage periods but not between PM1 and PM2 storage periods. The PPI column indicates significant results ($p < 0.1$) from the PPI analyses after performing small-volume corrections based on 10-mm-radius spheres centered on the relevant peak.

analyses, two of the four peak regions identified in the cue decoding analysis also showed enhanced coupling with RLPFC during PM intention storage, as did three of the five regions identified in the response decoding analysis, at least at a marginally significant threshold ($p < 0.1$) (Table 1). The probability of five or more of nine tests yielding at least a marginally significant effect, under the null hypothesis, is extremely low (binomial test: $p = 0.0009$). To examine the specificity of this effect for the regions identified in the MVPA analyses, PPI results were also examined for the five peak regions identified in the univariate contrast of PM store > baseline store, none of which was included in the PPI seed region mask. Only one of these five showed a significant effect at this threshold, not significantly different from results expected by chance, given the number of tests performed (binomial test, $p = 0.41$). Thus, functional connectivity appeared to be enhanced particularly with content-representing brain regions, rather than all brain regions involved in the task.

These results are consistent with the hypothesis that RLPFC supports prospective remembering via its interactions with distributed content-representing posterior brain regions. One prediction derived from this model might be that trials with relatively high RLPFC activation should be accompanied by relatively high intention decoding accuracy from posterior brain regions, because on these trials RLPFC should support a relatively strong representation of delayed intentions in these brain areas. Accuracy of intention decoding from posterior brain regions was therefore compared between trials with relatively high versus low RLPFC activity. This analysis was performed for all nine peak regions identified in the cue and response decoding analyses. Response decoding accuracy was significantly greater in left putamen on trials with relatively high versus low RLPFC activity ($p = 0.0036$; this level of statistical significance passes a Bonferroni correction for multiple comparisons across nine regions). None of the other regions identified in the decoding analyses showed a similar effect ($p > 0.15$, uncorrected). This result would be consistent with a causal role of RLPFC in maintaining representations of future intentions in the putamen. Seeing as this effect was not observed in any of the other eight regions identified in the cue and response decoding analyses, it is likely that it reflects specific interactions between RLPFC and putamen, rather than a global factor that differed between high and low

RLPFC activation trials (which would be expected to affect all nine brain regions equally).

Discussion

The present study demonstrates (1) increased RLPFC activation associated with storage (but not encoding or retrieval) of delayed intentions; (2) representation of the content of those intentions in distributed posterior brain regions, but not RLPFC itself; (3) increased coupling between RLPFC and content-representing posterior brain regions during storage of delayed intentions; and (4) enhanced discrimination of the content of intentions from one such region (left putamen) on trials with relatively high RLPFC activation. Together, these findings point toward a role of RLPFC in maintaining representations of forthcoming cues and appropriate responses stored in posterior brain regions, to enable realization of delayed intentions at the appropriate time. In other words, the present results support a retrieval mode rather than a goal stack model of the role of RLPFC in maintaining delayed intentions.

It is notable that some regions involved in processing current perceptual input, such as medial occipital cortex, also maintained a representation of the forthcoming cue for action. Likewise, parts of the motor system maintained a representation of the appropriate forthcoming response. These findings are consistent with other evidence showing that cortical regions specialized for processing particular types of low-level perceptual and motor information can also play content-specific mnemonic functions (Chieffi et al., 1999; Eschen et al., 2007; Harrison and Tong, 2009). It has previously been reported that representations of future intentions (which of two cognitive operations: addition vs subtraction) can be decoded from medial prefrontal cortex (Haynes et al., 2007). This study differed from the present paradigm in that decoding was from a period of the task where participants attended to their intended action, just before executing it. By contrast, the present study focused on the distinct scenario in which participants must maintain a delayed intention while they are occupied with an unrelated ongoing task. The region of medial prefrontal cortex identified by Haynes et al. (2007) was considerably posterior to the region identified here, in a part of the brain known to exhibit considerable fine-grained functional specialization (Gilbert et al., 2006, 2010).

Consistent with its role in representing a forthcoming cue for action, the anterior medial prefrontal region implicated in the present study has also been implicated in predictive coding of expected visual stimuli (Summerfield et al., 2006). A similar region was also identified by Soon et al. (2008), from which it was possible to decode which of two motor responses participants were preparing to make. By contrast, the more posterior region identified by Haynes et al. (2007) may be involved in higher-level representation of cognitive operations (i.e., task sets) (Monsell, 2003) rather than specific perceptual or motor representations. In this way, rather than intentions being represented in just one brain area, multiple distinct regions are proposed to represent different dimensions of our future intentions: perceptual triggers for action, specific responses, and higher-level task sets or representations of action schemas (Wood and Grafman, 2003). Meanwhile, RLPFC appears to play a supporting role, marshalling these posterior representations so that relevant cues are detected, ongoing behavior interrupted, and intended actions executed.

The idea that RLPFC plays a “content-free” role in PM is congruent with repeated findings of RLPFC activation in PM paradigms using a wide variety of different stimulus types and task demands (Okuda et al., 1998; Burgess et al., 2001, 2003; Simons et al., 2006; Gilbert et al., 2009). Furthermore, the present

results fit well with clinical investigations of PM in patients with frontal lobe lesions. Such patients have been thought to have a difficulty in bringing to mind delayed intentions at the appropriate time, rather than forming a representation of them (Burgess et al., 2000; Uretzky and Gilboa, 2010). In addition, some patients derive considerable benefit from nonspecific cues (e.g., an auditory beep) that remind them of the existence of a delayed intention but not its content (Fish et al., 2007). This would be expected if PM failures in such patients do result not from an inability to store the content of future intentions, but rather from an inability to activate the representations of such intentions in posterior brain regions. Of course, just because the present study could not decode particular aspects of the representational content of future intentions from RLPFC, this does not preclude the existence of any such representations within this brain region. Indeed, an interesting question for future research will be to ask which aspects of delayed intentions, if any, are represented in the spatial distribution of activation within RLPFC.

References

- Adda CC, Castro LH, Além-Mar e Silva LC, de Manreza ML, Kashiara R (2008) Prospective memory and mesial temporal epilepsy associated with hippocampal sclerosis. *Neuropsychologia* 46:1954–1964.
- Badre D, D'Esposito M (2009) Is the rostro-caudal axis of the frontal lobe hierarchical? *Nat Rev Neurosci* 10:315–326.
- Brandimonte G, Einstein G, McDaniel M (1996) Prospective memory: theory and applications. Hillsdale, NJ: Erlbaum.
- Braver TS, Bongiolatti SR (2002) The role of frontopolar cortex in subgoal processing during working memory. *Neuroimage* 15:523–536.
- Burgess PW, Veitch E, de Lacy Costello A, Shallice T (2000) The cognitive and neuroanatomical correlates of multitasking. *Neuropsychologia* 38:848–863.
- Burgess PW, Quayle A, Frith CD (2001) Brain regions involved in prospective memory as determined by positron emission tomography. *Neuropsychologia* 39:545–555.
- Burgess PW, Scott SK, Frith CD (2003) The role of the rostral frontal cortex (area 10) in prospective memory: a lateral versus medial dissociation. *Neuropsychologia* 41:906–918.
- Chieffi S, Allport DA, Woodin M (1999) Hand-centered coding of target location in visuo-spatial working memory. *Neuropsychologia* 37:495–502.
- Christoff K, Keramatian K, Gordon AM, Smith R, Mädler B (2009) Prefrontal organization of cognitive control according to levels of abstraction. *Brain Res* 1286:94–105.
- Eichenbaum H, Cohen NJ (2001) From conditioning to conscious recollection: memory systems of the brain. New York: Oxford UP.
- Einstein GO, McDaniel MA, Williford CL, Pagan JL, Dismukes RK (2003) Forgetting of intentions in demanding situations is rapid. *J Exp Psychol Appl* 9:147–162.
- Einstein GO, McDaniel MA, Thomas R, Mayfield S, Shank H, Morrisette N, Breneiser J (2005) Multiple processes in prospective memory retrieval: factors determining monitoring versus spontaneous retrieval. *J Exp Psychol Gen* 134:327–342.
- Eschen A, Freeman J, Dietrich T, Martin M, Ellis J, Martin E, Kliegel M (2007) Motor brain regions are involved in the encoding of delayed intentions: a fMRI study. *Int J Psychophysiol* 64:259–268.
- Fish J, Evans JJ, Nimmo M, Martin E, Kersel D, Bateman A, Wilson BA, Manly T (2007) Rehabilitation of executive dysfunction following brain injury: “content-free” cueing improves everyday prospective memory performance. *Neuropsychologia* 45:1318–1330.
- Friston KJ, Holmes AP, Worsley KJ, Poline J-P, Frith CD, Frackowiak RSJ (1995) Statistical parametric maps in functional imaging: a general linear approach. *Hum Brain Mapp* 2:189–210.
- Friston KJ, Buechel C, Fink GR, Morris J, Rolls E, Dolan RJ (1997) Psychophysiological and modulatory interactions in neuroimaging. *Neuroimage* 6:218–229.
- Friston KJ, Glaser DE, Henson RN, Kiebel S, Phillips C, Ashburner J (2002) Classical and Bayesian inference in neuroimaging: applications. *Neuroimage* 16:484–512.
- Gilbert SJ, Spengler S, Simons JS, Steele JD, Lawrie SM, Frith CD, Burgess PW (2006) Functional specialization within rostral prefrontal cortex (area 10): a meta-analysis. *J Cogn Neurosci* 18:932–948.
- Gilbert SJ, Gollwitzer PM, Cohen AL, Burgess PW, Oettingen G (2009) Separable brain systems supporting cued versus self-initiated realization of delayed intentions. *J Exp Psychol Learn Mem Cogn* 35:905–915.
- Gilbert SJ, Henson RN, Simons JS (2010) The scale of functional specialization within human prefrontal cortex. *J Neurosci* 30:1233–1237.
- Gollwitzer PM (1999) Implementation intentions: strong effects of simple plans. *Am Psychol* 54:493–503.
- Harrison SA, Tong F (2009) Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458:632–635.
- Haynes JD, Sakai K, Rees G, Gilbert S, Frith C, Passingham RE (2007) Reading hidden intentions in the human brain. *Curr Biol* 17:323–328.
- Kliegel M, McDaniel MA, Einstein GO (2008) Prospective memory: cognitive, neuroscience, developmental, and applied perspectives. Mahwah, NJ: Erlbaum.
- Koechlin E, Basso G, Pietrini P, Panzer S, Grafman J (1999) The role of the anterior prefrontal cortex in human cognition. *Nature* 13:148–151.
- Koechlin E, Ody C, Kouneiher F (2003) The architecture of cognitive control in the human prefrontal cortex. *Science* 302:1181–1185.
- Kriegeskorte N, Goebel R, Bandettini P (2006) Information-based functional brain mapping. *Proc Natl Acad Sci U S A* 103:3863–3868.
- McDaniel MA, Einstein GO (2000) Strategic and automatic processes in prospective memory retrieval: a multiprocess framework. *Appl Cogn Psychol* 14:S127–S144.
- Monsell S (2003) Task switching. *Trends Cogn Sci* 7:134–140.
- Okuda J, Fujii T, Yamadori A, Kawashima R, Tsukiura T, Fukatsu R, Suzuki K, Ito M, Fukuda H (1998) Participation of the prefrontal cortices in prospective memory: evidence from a PET study in humans. *Neurosci Lett* 253:127–130.
- Reynolds JR, West R, Braver T (2009) Distinct neural circuits support transient and sustained processes in prospective memory and working memory. *Cereb Cortex* 19:1208–1221.
- Rowe JB, Toni I, Josephs O, Frackowiak RS, Passingham RE (2000) The prefrontal cortex: response selection or maintenance within working memory? *Science* 288:1656–1660.
- Rowe JB, Sakai K, Lund TE, Ramsoy T, Christensen MS, Baare WF, Paulson OB, Passingham RE (2007) Is the prefrontal cortex necessary for establishing cognitive sets? *J Neurosci* 27:13303–13310.
- Sakai K, Passingham RE (2003) Prefrontal interactions reflect future task operations. *Nat Neurosci* 6:75–81.
- Simons JS, Spiers HJ (2004) Prefrontal and medial temporal lobe interactions in long-term memory. *Nat Rev Neurosci* 4:637–648.
- Simons JS, Schölvinck ML, Gilbert SJ, Frith CD, Burgess PW (2006) Differential components of prospective memory? Evidence from fMRI. *Neuropsychologia* 44:1388–1397.
- Smith RE (2003) The cost of remembering to remember in event-based prospective memory: investigating the capacity demands of delayed intention performance. *J Exp Psychol Learn Mem Cogn* 29:347–361.
- Soon CS, Brass M, Heinze HJ, Haynes JD (2008) Unconscious determinants of free decisions in the human brain. *Nat Neurosci* 11:543–545.
- Squire LR, Stark CE, Clark RE (2004) The medial temporal lobe. *Annu Rev Neurosci* 27:279–306.
- Summerfield C, Egner T, Greene M, Koechlin E, Mangels J, Hirsch J (2006) Predictive codes for forthcoming perception in the frontal cortex. *Science* 314:1311–1314.
- Tulving E (1983) Elements of episodic memory. Oxford: Clarendon.
- Uretzky S, Gilboa A (2010) Knowing your lines but missing your cue: rostral prefrontal lesions impair prospective memory cue detection, but not action-intention superiority. *J Cogn Neurosci* 22:2745–2757.
- Visscher KM, Miezin FM, Kelly JE, Buckner RL, Donaldson DI, McAvoy MP, Bhalodia VM, Petersen SE (2003) Mixed blocked/event-related designs separate transient and sustained activity in fMRI. *Neuroimage* 19:1694–1708.
- West R (2008) The cognitive neuroscience of prospective memory. In: Prospective memory: cognitive, neuroscience, developmental, and applied perspectives (Kliegel M, McDaniel MA, Einstein GO, eds), pp 261–282. Mahwah, NJ: Erlbaum.
- Wood JN, Grafman J (2003) Human prefrontal cortex: processing and representational perspectives. *Nat Rev Neurosci* 4:139–147.