Retrieval Goal Modulates Memory for Context

Rachael L. Elward and Michael D. Rugg

Abstract

■ fMRI was employed to assess whether the contents of recollection vary according to retrieval goal. At study, visually presented words were superimposed on urban or rural scenes or a gray background. The word–background pairs were presented in one of three spatial locations. During a scanned test phase, studied and unstudied words were presented. Two different source memory tasks were randomly interleaved. In the "background" task, the requirement was to judge whether the word had been presented against one of the two classes of scene, as opposed to the alternate class or the gray background. In the "location" task, discrimination was between words presented in one of the two lateral locations and words presented in either of the alternate locations. In both tasks, unstudied words required a separate response. In the background task, words studied against scenes elicited greater activity in parahippocampal and retrosplenial cortex than did words studied against the gray background, consistent with prior reports of scene reinstatement effects. Reinstatement effects were also evident in the location task. Relative to the background task, however, the effects were attenuated in parahippocampal cortex. In other regions, including medial prefrontal and posterior cingulate cortex, activity elicited in the location task by items associated with scenes was lower than that elicited by items presented on the gray background. The findings are interpreted as evidence that contextual retrieval is partially modulated by retrieval goal. ■

INTRODUCTION

An important component of episodic retrieval (recollection) is the selection of those aspects of a memory representation that are relevant to the goal of the retrieval attempt. For example, if asked about the weather during your drive to work earlier in the day, remembering in addition the music that was playing on the car radio would necessitate engagement of resource- and time-consuming "postretrieval monitoring" operations to select the features of the retrieved information relevant to the question (cf. Halamish, Goldsmith, & Jacoby, 2012). Thus, the ability to control which features of a memory representation are brought to mind would be highly beneficial. Previous research has led to the identification of control processes that facilitate retrieval of goal-relevant episodes and reduce the likelihood of retrieving goal-irrelevant episodes, with a concomitant reduction of demands on postretrieval processing (e.g., Herron & Wilding, 2006; Jacoby, Shimizu, Daniels, & Rhodes, 2005; Rugg, 2004; Rugg & Wilding, 2000). It is currently uncertain, however, whether goal-relevant features of a single episodic memory representation can be retrieved at the expense of goal-irrelevant features.

One possibility is that retrieval processing can indeed be controlled so as to prevent the retrieval of irrelevant information. For instance, when asked how the weather was when you drove to work, memory search could be biased to favor retrieval of weather-related information and to minimize the likelihood of retrieving other, irrelevant features. Thus, different retrieval goals might lead to the differential processing of a retrieval cue in much the same way as has been proposed to occur for entire episodes (Herron & Wilding, 2006; Jacoby et al., 2005; Rugg & Wilding, 2000; see also Anderson & Bjork, 1994). Alternately, selective retrieval might be accomplished not through a biased search process, but by a "gating" operation that allows goal-relevant information contained within a memory representation to be reinstated (see below) while suppressing reinstatement of irrelevant information (cf. Wimber, Alink, Charest, Kriegeskorte, & Anderson, 2015). This latter mechanism is arguably more compatible than the biased search account with proposals that the putative hippocampally mediated "pattern completion" process that causes a stored memory representation to become active is automatic (e.g., Halamish et al., 2012) and, typically, "all-or-none" (i.e., the representation is either activated in its entirety or not all; Norman & O'Reilly, 2003).

The alternative possibility, of course, is that episodic retrieval is not modulated by retrieval goal. By this account, all accessible features of an event are retrieved regardless of the goal of the retrieval attempt. The selection of goal-relevant features of a memory representation thus depends on postretrieval processes that operate after both goal-relevant and goal-irrelevant episodic information has been reinstated. This view has been advocated by Kuhl, Johnson, and Chun (2013; see below).

University of Texas at Dallas

Journal of Cognitive Neuroscience 27:12, pp. 2529–2540
doi:10.1162/jocn_a_00878

One way to decide between these different possibilities is by exploiting the phenomenon of cortical reinstatement. It is well established that episodic retrieval is associated with reinstatement in the brain of some of the processes and representations that were active when the episode was initially experienced (for reviews, see Rugg, Johnson, & Uncapher, 2015; Danker & Anderson, 2011). For example, cortical regions engaged during the processing of visual scene information, such as the retrosplenial cortex (RSC) and the parahippocampal cortex (PHC), are active not only when the information is processed online but also when it is subsequently recollected (Gordon, Rissman, Kiani, & Wagner, 2014; Johnson, Suzuki, & Rugg, 2013; Kuhl, Rissman, & Wagner, 2012; Staresina, Henson, Kriegeskorte, & Alink, 2012; Johnson & Rugg, 2007). Thus, if retrieval of episodic information can indeed be prevented or, at least, attenuated when it is not relevant to the retrieval goal, then its cortical reinstatement should be weaker relative to when the information is goal relevant. Alternatively, if goal-relevant information is only selected postretrieval, reinstatement effects should not be modulated by the retrieval goal.

The present experiment was designed to adjudicate between these alternatives. During the study phase, words were presented against one of two classes of background context and in one of three spatial locations. One context class comprised informationally rich visuospatial scenes, whereas the alternate class was a uniform gray background. During the test phase, we varied, trial-by-trial, whether the retrieval task required retrieval of a studied test word's study background or its study location. We predicted that when background information was relevant to the retrieval task scene reinstatement effects would be evident in regions—such as RSC and PHC—engaged during online scene processing, replicating prior findings (see above). The key question concerned the magnitude and extent of "scene reinstatement" effects when location, rather than background, information was the task-relevant feature. According to the first possibility outlined above, the effects should be much reduced, whereas according to the second possibility, the effects should be unaffected by the change in retrieval goal.

A prior study, employing a rationale very similar to the one just articulated, also addressed the question of whether reinstatement effects are modulated by retrieval goal. Kuhl et al. (2013) used multivoxel pattern analysis (MVPA) to examine whether reinstatement of scene and face information was stronger when the retrieval task required a judgment about whether a word had been studied with one or other of these classes of contextual feature than when a judgment about study location was required. The authors found no evidence that reinstatement effects in cortical regions selectively engaged by face and scene stimuli (fusiform cortex and PHC) were modulated by retrieval goal. Rather, goal-sensitive effects were identified in prefrontal and dorsal parietal cortex. These findings were interpreted as evidence that selection of goal-relevant mnemonic information takes place postretrieval, that is, as evidence favoring the second of the two possibilities outlined above. It remains to be established how well this finding generalizes to other experimental settings. Notably, the study of Kuhl and colleagues employed repeated study–test cycles that contained only 12 study trials and 8 test trials per cycle. The use of such short lists, along with the brief study–test delay that was imposed, may have made the encoded representations of the study trials highly accessible, encouraging a "postretrieval" strategy (as was reported by participants on a postexperiment questionnaire). By contrast, this study employed a single study–test cycle and more demanding test tasks, conditions that arguably are more conducive to selective retrieval.

METHODS

Participants

Twenty-two right-handed, English-speaking adults, aged 18–29 years, were recruited from the local community and from the student and staff bodies of the University of Texas at Dallas. All participants indicated by self-report that they were free from neurological and psychiatric disorders and gave informed consent to participate in the experiment. They received $30 per hour as compensation for their time. Two participants were excluded from all analyses because of excessive false alarm rates (rates of 0.70 and 0.76, respectively). The remaining 20 participants were included in the fMRI analyses that were based on correctly recognized items (see below). For the analyses based on items attracting accurate source memory judgments, a further two participants were excluded because of low trial numbers for one or more of the events of interest (fewer than five trials). Thus, 18 participants contributed data to these analyses. Behavioral data are presented from the 18 participants included in both sets of analyses (the results were unchanged when the data from all 20 participants were included). The experiment was approved by the institutional review boards of the University of Texas, Southwestern Medical Center, and the University of Texas at Dallas.

Experimental Items

Study items were visually presented words superimposed on backgrounds that were depictions of a rural scene, an urban scene, or a gray square. The scenes, which were trial-unique, were selected from the Computational Visual Cognition Laboratory database (cvcl.mit.edu/database.htm). The rural scenes were images of open countryside, whereas the urban scenes comprised images of city streets. Images that had a salient object or animal in the foreground that closely resembled other images in the item pool or that contained legible words were not used. Study words were selected from the MRC psycholinguistic database (websites.psychology.uwa.edu.au/school/MRCDatabase/uwa_mrc.htm) and were restricted to concrete nouns ranging between 3 and 10 characters in length. The words were allocated to left, right, and center screen locations and to the rural, urban, and gray backgrounds, such that there were 20 study trials belonging to each of the nine possible combinations of location and background type. Words were rotated across participants so that each word was equally likely to appear at each location and against each class of background image or to appear as a new test item. One hundred eighty critical study trials were presented in total, as well as six filler items (two at the start of each block), which were not included in any of the reported analyses.

Experimental Procedure

A schematic depiction of the experimental procedure is given in Figure 1. The encoding phase took place in a testing room adjacent to the MRI scanner. Word–background pairs were presented on a laptop computer, and participants responded on the keyboard. Each trial began with the presentation of a black fixation cross for 200 msec. The study word was then presented in red font, with the background image appearing 200 msec after word onset. Participants were instructed to imagine the object denoted by the word moving around inside the scene or gray square and to rate the pleasantness of their mental image on a scale ranging from 1 (*unpleasant*) to 3 (*pleasant*). The word–background pair remained on the screen for 5550 msec, during which time participants made their response. A 200 msec ISI followed, during which a black fixation cross was presented. A short practice sequence was administered before the study phase proper. The study phase was divided into two blocks, each with 90 critical trials preceded by two filler trials. A brief rest was provided between the blocks.

Participants were not informed about the nature of the retrieval tasks until after the study phase. Two source memory tasks were randomly interleaved. In the "background" task, participants were required to recall the class of background that was presented with each studied word during the study phase. This information was not relevant, however, in the location task: Here, participants were required to recall the location at which a studied item had been presented at study. For each participant, one location (left or right) and one scene category (rural or urban) were assigned as the "target" for the duration of the memory test. Target locations and scene categories were counterbalanced across participants, such that each location and scene category was assigned as the target and the nontarget an equal number of times. Each test trial began with a task cue that corresponded to either the target location (e.g., "Left?") or to the target scene category (e.g., "Rural?"). The cue signaled the retrieval task to be performed on the upcoming trial. For example, when "Rural" was the target background, each trial in the background task began with the cue "Rural?" Participants were instructed to interpret the cue as, "was the upcoming word studied with a rural scene?" and to answer with one of four responses: "yes" (indicating recollection of a target scene), "no" (indicating recollection of a nontarget scene or a gray context), "don't know" (indicating the item was old but that the background context information could not be recollected), or "new" (indicating that the test word was unstudied). Analogously, a location cue (e.g., "Right?") signaled the question "Was the upcoming word studied on the right side of the screen?" Responses were

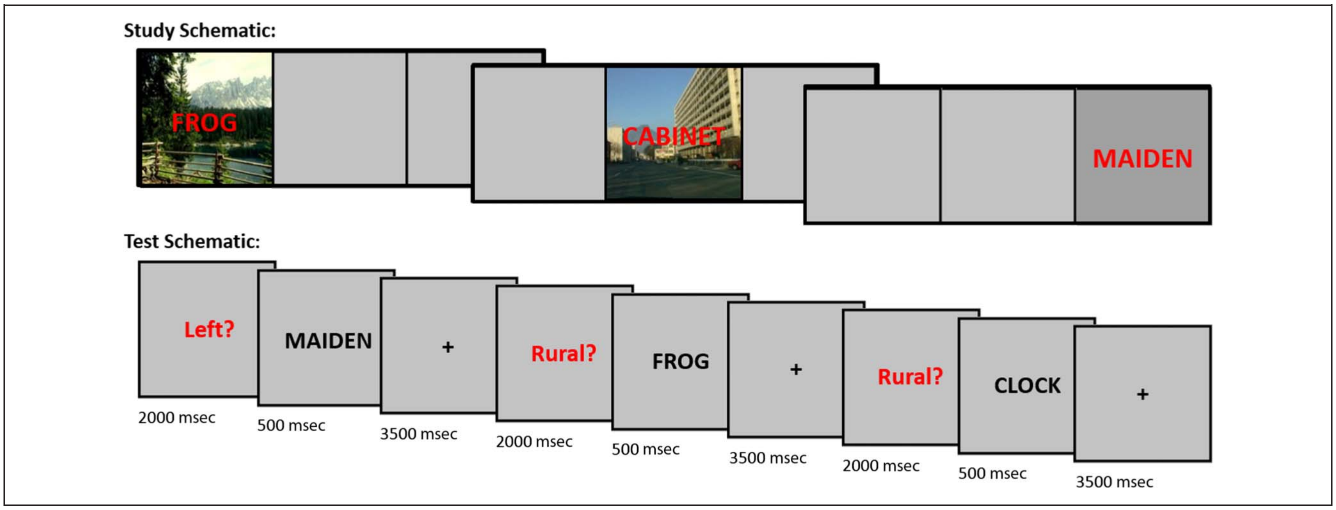

Figure 1. Schematic representation of the study and test phases. At study (top), boxes were displayed on the left, right, and center of the screen. On each trial, a word and background image were presented in one of the three boxes. At test (bottom), two tasks were interleaved. The task was indicated by a cue word presented in red at the onset of each trial.

"yes" (indicating that the item was presented at the target location), "no" (indicating that the item was presented at the nontarget location or centrally), "don't know" (indicating the item was old but that location information was not recollected), or "new." For both tasks, participants were instructed to respond "yes" or "no" only when confident of their judgment. They were instructed that if they were unsure about their judgment, then they should press the "don't know" key. Participants were also instructed that if they were unsure whether a word had been studied, then they should respond on the "new" key.

The test instructions were administered outside the scanner and were followed by a short practice test, which included all of the words from the practice study phase as well as new (unstudied) words. The practice was repeated inside the scanner while survey and reference scans were acquired. During the test proper, all of the words presented in the study phase were re-presented, interspersed with 60 new words.

On each test trial, a task cue was presented for 2000 msec, the test word (retrieval cue) was presented for 500 msec, and a fixation cross was then presented for 3500 msec, during which period participants made their response. The screen was then blanked for 100 msec before the next task cue was presented. The test phase was divided into three runs with a short break between each run. During the breaks, participants were reminded of the task instructions and response options. Each test run included 60 randomly interspersed null trials, when only a fixation cross was presented. Participants were instructed to maintain fixation during these trials and await the upcoming task cue.

MRI Data Acquisition

BOLD T2*-weighted echoplanar functional images (SENSE factor 1.5, flip angle 70°, 80 × 80 matrix, field of view = 24 cm, repetition time = 2000 msec, echo time = 30 msec) and T1-weighted anatomical images (MPRAGE sequence, 240 × 240 matrix, 1 mm isotropic voxels) were acquired with a 3-T Philips Achieva MRI (Philips Medical Systems, Andover, MA) scanner equipped with a 32-channel receiver head coil. Three hundred and twenty-six functional volumes were acquired during each of the three test runs. Each volume comprised 34 slices, acquired in an ascending sequence, oriented parallel to the AC–PC line (3 mm isotropic voxels, 1 mm interslice gap). The first five volumes of each scanning session were discarded to allow equilibration of tissue magnetization. For each session, mean signal intensity across volumes and voxels was adjusted to a nominal value of 100.

fMRI Data Analysis

Statistical Parametric Mapping (SPM8, Wellcome Department of Cognitive Neurology, London, United Kingdom), run under Matlab R2012a (The MathWorks, Inc., Natick, MA) was used for fMRI data analysis. Functional images were subjected to realignment (to the mean image), slice timing correction (using the 17th slice as the reference), reorientation, spatial normalization to a standard EPI template (based on the Montreal Neurological Institute [MNI] reference brain; Cocosco, Kollokian, Kwan, & Evans, 1997) and smoothing with an 8 mm FWHM Gaussian kernel. Each participant's structural volume was normalized to the MNI T1 template before averaging to create an across-participant mean image. Functional analysis was performed using a general linear model (GLM) in which a delta function was used to model neural activity at test item onset. The function was convolved with a canonical hemodynamic response function to model the predicted BOLD response (Friston et al., 1995).

Our analysis strategy required us to construct two pairs of first- and second-level GLMs. Model 1 was employed to identify "scene reinstatement" effects, that is, regions that showed greater activity in response to test words that were previously associated with a scene than with a gray background. Once these regions were identified, they were interrogated to compare the magnitude of scene reinstatement effects across the two tasks. All correctly recognized items, regardless of the accuracy of the associated source judgment, were included in these analyses. By doing this, we avoided the bias in favor of the background task that would have existed had the analysis been limited to items attracting correct source judgments. The bias would have arisen because a correct source judgment in the background task required successful retrieval of the test item's studied background, whereas correct source judgments in the location task did not depend on successful recollection background information. Hence, background information may not have been retrieved each time a correct location judgment was made. Thus, a comparison restricted to correct source judgments would have been biased toward a finding of greater scene reinstatement effects in the background task regardless of any influence of retrieval goal. If however, consistent with the null hypothesis, the probability of successful retrieval of background context was equivalent in the two tasks, then the proportion of correctly recognized test items (collapsed across accurate and inaccurate source judgments) associated with retrieval of background information would also be equivalent and the neural correlates of scene retrieval would not be expected to differ.

By contrast, Model 2 was employed to examine the activity elicited by correctly recognized items that were also accorded a correct source judgment. Here, the aim was to identify regions where activity varied as a function of task and background context when the task-relevant source feature was successfully retrieved. We were particularly interested in whether it was possible to identify regions that were differentially active when participants succeeded in making a correct location judgment for items studied in association with scenes rather than gray backgrounds (see below). In short, in this analysis, we sought to identify regions that might support the cognitive processes engaged to control the recollection of competing contextual features.

The first-level GLMs of both of the foregoing models employed six regressors representing motion-related variance (three for rigid body translation and three for rotation) as well as regressors modeling the separate scan sessions and the across-scan mean. An AR(1) model was used to estimate and correct for nonsphericity of the error covariance (Friston et al., 2002) in both models.

Model 1 included nine trial types at the first level: (1) correct recognition in the background task of items studied with a target scene, (2) correct recognition in the background task of items studied with a nontarget scene, (3) correct recognition in the background task of items studied with a gray background, (4) correct rejection of a new item in the background task, (5) correct recognition in the location task of items studied with scenes in the target location, (6) correct recognition in the location task of items studied with gray background in the target location, (7) correct recognition in the location task of items studied with scenes in the nontarget location, (8) correct recognition in the location task of items studied with a gray background in the nontarget location, and (9) correct rejection of a new item in the location task. An additional event type of no interest included false alarms (new items incorrectly endorsed as old), item misses (old items incorrectly judged new), and trials associated with multiple or omitted responses. Four trial types were carried forward to the second level. These comprised the trials associated with making a correct "old" response to an item paired with a nontarget scene or a gray background item in each task (Trial types 2, 3, 7, and 8 above). The parameter estimates associated with these conditions were entered into a whole-brain second-level ANOVA model (factors of Task [background task vs. location task] and Background context [scene context vs. gray context]).

Model 2 was similar to Model 1 except that source correct judgments were separated from incorrect and source Don't Know (DK) judgments. The model contained nine trial types: (1) source correct responses in the background task for items studied with a target scene, (2) source correct responses in the background task for items studied with a nontarget scene, (3) source correct responses in the background task for items studied with a gray background, (4) source incorrect/source DK responses in the background task, (5) correct rejections in the background task, (6) source correct responses in the location task for items studied with scenes, (7) source correct responses in the location task for items studied with gray backgrounds, (8) source incorrect/source DK responses in the location task, (9) correct rejections in the location task. As in the prior model, an additional event type included trials of no interest (e.g., false alarms, misses, and omitted responses). As already noted, 18 participants had sufficient (five or more) trial numbers in each of these trial types to be included in this model. Parameter estimates corresponding to all "source correct" responses were carried forward to the second-level GLM, with the exception of Trial type 1 (source correct responses in the background task to items studied with a target scene); this trial type was not carried forward because no contrasts of interest included this condition. The parameter estimates for the four relevant conditions (corresponding to Trial types 2, 3, 6, and 7 above) were entered into a whole-brain ANOVA structured as described previously.

Note that trials in the background task that were associated with the target scene (Trial type 1 in both models) were excluded from both of the foregoing fMRI analyses. The rationale was to avoid the response confound in this task when contrasting items paired with scenes or gray backgrounds that would have come about because "targets" trials were always associated with a studied scene. This meant that correct responses to items paired against gray backgrounds in the background task were always associated with a "no" response, whereas correct responses to scenes were divided between "yes" and "no" responses. The confound is obviated by excluding target items in the background task and comparing only those test trials that were associated with a gray background with a nontarget scene.

F and *t* contrasts derived from each of the second-level models were height thresholded at $p < .001$ (uncorrected) and combined with a 21-voxel cluster extent threshold, giving a corrected whole-brain cluster-wise significance level of $p < .05$ as estimated using Monte Carlo simulations implemented in AFNI (afni.nimh.nih.gov/afni). Coordinates of significant effects are reported in MNI space. Effects of interest are displayed on sections of the across-participant mean normalized structural image.

RESULTS

Behavioral Data

Memory Performance

Hit rates in the background and location tasks were 0.93 and 0.92, respectively, against false alarm rates of 0.13 and 0.18. Item memory (indexed as hit rate – false alarm rate, or pR) was significantly higher for the background task (mean = 0.80, *SD* = 0.10) than for the location task (mean = 0.74, *SD* = 0.15), $t(17) = 2.44$, $p < .05$. Within each task, hit rates were contrasted according to study background (target scene, nontarget scene, and gray background; see Figure 2, left). In neither task did the hit rates differ significantly (max $F = 1.42$).

Source Memory Accuracy

Source recollection—"pSr"—was assessed using an index derived from a single high threshold model of memory

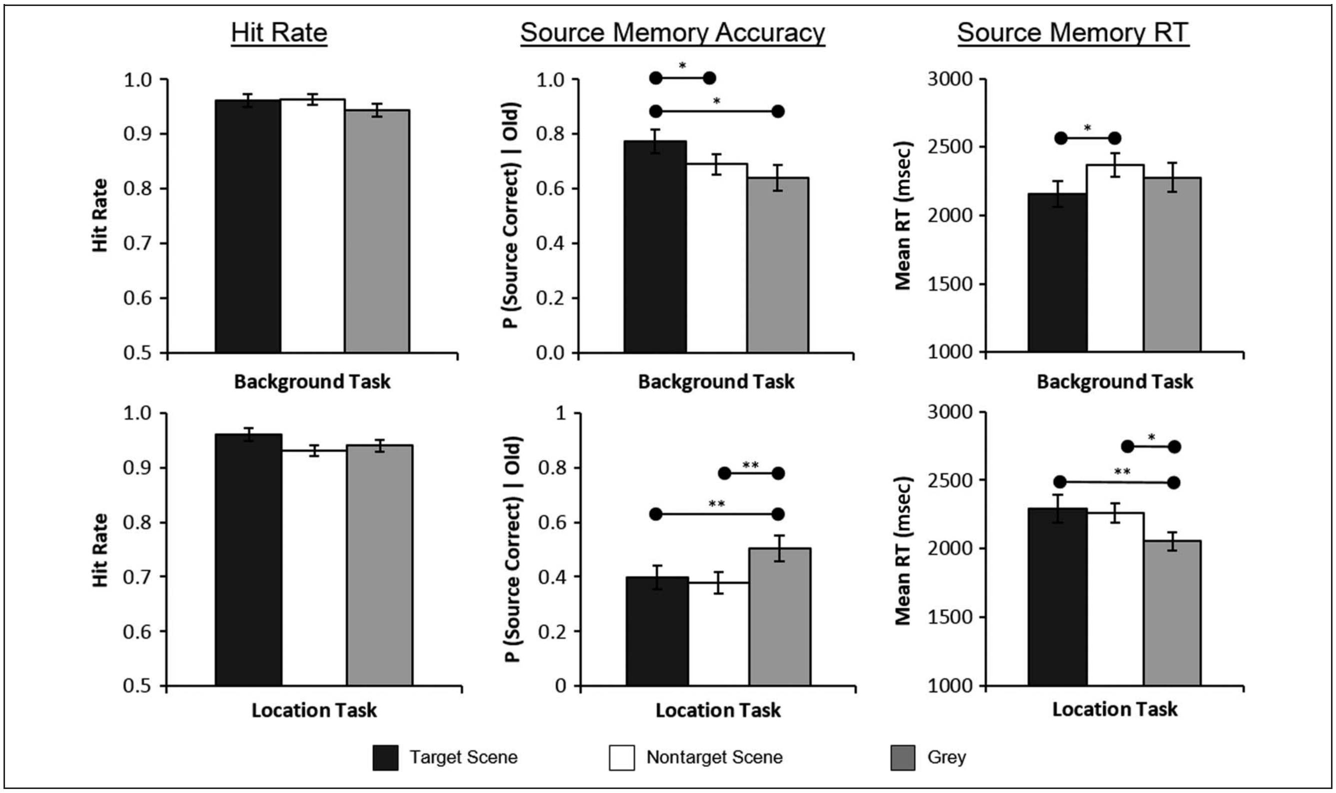

Figure 2. Behavioral data from the test phase. Error bars indicate ±*SEM*. $^*p < .05$, $^{**}p < .01$.

judgments (Snodgrass & Corwin, 1988; for a previous application, see Gottlieb et al., 2010). The index corrects for guessing using the formula, pSr = [p(source hit) − 0.5(1 − p(source don't know)]/[1 − 0.5(1 − p(source don't know))]. Thus, an index of 1 indicates perfect performance, whereas 0 indicates chance performance. Mean pSr was 0.53 (*SD* = 0.21) and 0.19 (*SD* = 0.13) in the background and location tasks, respectively; these means differed significantly ($t(17) = 8.06$, $p < .001$). One-sample *t* tests indicated that source memory performance exceeded chance in both tasks (minimum $t = 6.02$, $p < .001$).

Source Accuracy for Each Context Type

Accuracy of source memory was contrasted across items paired with each type of background context (target scenes, nontarget scenes, and gray backgrounds) separately for each task (see Figure 2, center). Source accuracy was estimated for each class of background as the probability of a correct source judgment given that the item had been correctly judged as old. The data were analyzed with two repeated-measures ANOVAs, one for each task. In the Background task, there was a significant effect of background context ($F(1.7, 29.5) = 5.41$, $p < .05$). Follow-up *t* tests indicated that participants were more accurate when making "target" than "nontarget" responses ($p < .05$). There were no differences in accuracy between items associated at encoding with nontarget scenes or gray backgrounds.

For the purpose of the fMRI analyses, scenes designated as targets and nontargets were collapsed into a single "scene context" category in the location task, because the different scene types were not behaviorally relevant in this task. Nonetheless, here we compared source memory accuracy in the location task according to each class of background context at encoding (target scene, nontarget scene, or gray background) in the same way as was just described for the background task. The ANOVA revealed a main effect of Context ($F(1.8, 30.9) = 8.43$, $p < .001$). Follow-up *t* tests indicated that location judgments were more accurate for items studied against the gray background than against either class of scenes (*p*s < .01). There was no difference in source accuracy between words studied in association with a target or a nontarget scene.

RT

RT data for correct source judgments are presented in Figure 2 (right). For each retrieval task, these data were subjected to repeated-measures ANOVA with a single factor of Background context (target scene vs. nontarget scene vs. gray background). For the background task, there was a significant Context effect, $F(1.6, 26.9) = 3.81$, $p < .05$. Follow-up *t* tests revealed that participants were faster to endorse a target scene with a "yes" response than they were to reject a nontarget scene with a "no" response, $t(17) = 2.87, p < .05$. No other contrasts were significant. For the location task, there was again a significant effect of Background context ($F(1.9, 32.6) = 5.65, p < .01$). Follow-up t tests indicated that this effect was driven by shorter RTs to words associated at study with the gray context than with either class of scenes (ps $< .05$). There were no differences between words studied in association with target and nontarget scenes.

fMRI Results

As was noted above (see fMRI Data Analysis) we used two different approaches to analyze the fMRI findings, and accordingly, the findings are presented in two sections. First, we contrasted the magnitude of scene effects in each task. Second, we identified regions where activity was modulated by both task and background context.

Scene Reinstatement Effects according to Retrieval Task

We employed the directional main effect of the ANOVA (Scene > Gray) to identify, unbiased by task, where words studied with scenes elicited greater activity than words studied with gray backgrounds. The contrast identified a single 820 voxel cluster with peaks in left ($-24, -37, -17$, peak $Z = 5.79$) and right ($21, -37, -20$, $Z = 3.65$) PHC. The cluster extended into the RSC bilaterally (peaks at $-18, -55, 7$, $Z = 3.78$, and $12, -55, 7$, $Z = 5.33$). These effects are illustrated in Figure 3.

To characterize the effects according to task and region, parameter estimates were extracted and averaged from a 3-mm radius sphere centered on each of the aforementioned four peaks. The data were entered into two ANOVAs (one for PHC and one for RSC) with factors in each case of Task (background vs. location), Hemisphere, and Background context (scene vs. gray). For PHC, the ANOVA revealed a trend toward a three-way interaction ($F(1, 19) = 3.43, p = .08$). More importantly, there was a significant Task × Background context interaction ($F(1, 19) = 12.70$, $p < .005$), such that background context memory effects were smaller in the location task than in the background task (see Figure 3). When the data from each hemisphere were analyzed separately, a significant Task × Background context interaction was evident in the right hemisphere, $F(1, 19) = 13.78, p < .001$, whereas the corresponding effect in the left hemisphere was a trend only, $F(1, 19) = 3.89, p < .07$. Pairwise comparisons between scene and gray backgrounds were conducted for each task and hemisphere. In the background task, activity associated with words encoded in scene contexts was reliably greater than that for words encoded against the gray

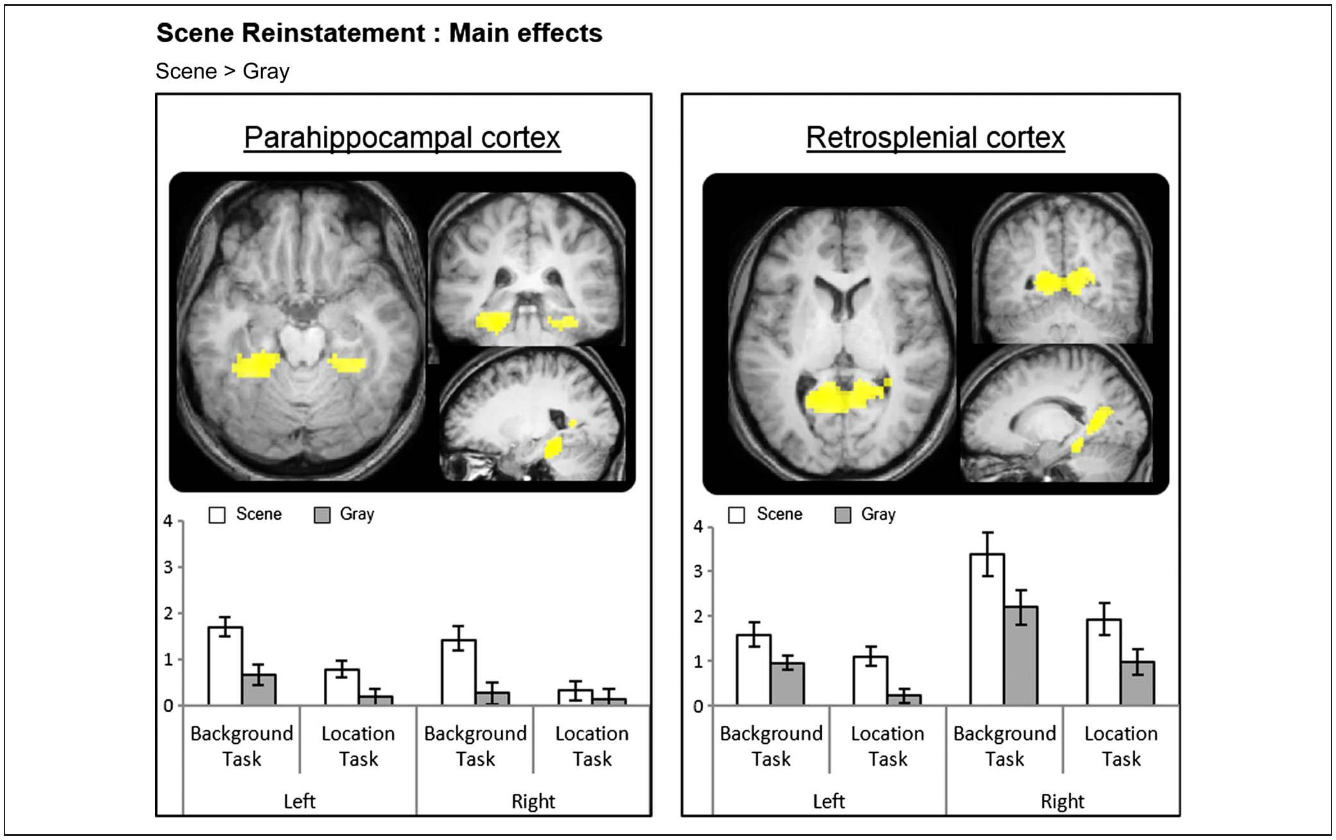

Figure 3. Top: Effects identified by the contrast between test words that had been studied against scene relative to gray backgrounds, collapsed across task and source accuracy. Bottom: Mean parameter estimates (arbitrary units) derived from a 3-mm sphere around the peak voxels in PHC and RSC in each hemisphere plotted as a function of task and background type. Error bars indicate ±1 *SEM*.

background (left $t(19) = 7.23, p < .001$); right $t(19) = 5.11$, $p < .001$). In the location task, the Scene > Gray effect was reliable in the left hemisphere, ($t(19) = 2.98, p < .01$), but no effect was evident on the right ($t < 1$).

In RSC, the ANOVA revealed a three-way interaction between Task, Hemisphere, and Background context ($F(1, 19) = 10.27, p < .005$). To characterize this interaction, the data from each hemisphere were analyzed separately. In the left hemisphere, the Task × Background context interaction was nonsignificant, ($F(1, 19) = 1.32$). There was a reliable main effect of task ($F(1, 19) = 19.36, p < .001$), however, such that the mean level of activity elicited in the background task was significantly greater than that in the location task. The pattern of effects in the right hemisphere was similar to that on the left; notably, there was again no evidence for a Task × Background context interaction ($F < 1$), but there was a main effect of Task ($F(1, 19) = 41.26, p < .001$) driven, as on the left, by greater activity in the background task than in the location task. Pairwise t tests revealed significant Scene > Gray context effects in both hemispheres and for each task (minimum $t(19) = 3.27, p < .01$).

To directly contrast the profiles of the Scene > Gray context effects in PHC and RSC, we conducted a third ANOVA that incorporated the factors of Region, Task, Hemisphere, and Background context. The four-way interaction was not significant ($F < 1$), but there was a significant interaction between Region, Task, and Background context ($F(1, 19) = 10.18, p < .005$), consistent with the foregoing findings indicating that Scene > Gray effects were modulated by task in PHC but not in RSC.

Finally, in a complementary analysis, we employed inclusive masking to identify voxels where reinstatement effects were modulated according to task. To achieve this, the directional main effect of the Scene > Gray contrast (thresholded at $p < .001$) was masked inclusively by the interaction between task (background vs. location) and background context (Scene vs. Gray), thresholded at $p < .05$, two-sided. As is evident from Figure 4, the outcome of this procedure was consistent with the ROI-based analyses described above. A cluster of 28 voxels in right PHC demonstrated a significant Task × Background context interaction (peak at 21, −34, −20, $Z = 3.16, p < .001$), and this was accompanied by a smaller cluster in left PHC (peak at −21, −34, −20, 11 voxels, $Z = 2.41, p < .01$). No voxels demonstrating a significant interaction could be identified in RSC.

Source Correct Responses

A whole-brain interaction contrast was performed to identify regions where the effects of background context were modulated by task when source judgments were accurate (Figure 5). The contrast identified five clusters: two of these were adjacent to one another in medial pFC (mPFC), one was localized to the posterior cingulate cortex (PC), another to the left medial-temporal lobe, and a final cluster was located in the left temporal pole (see Table 1). For each task separately, follow-up pairwise tests were conducted to contrast the activity elicited by items studied with each type of background context (scene vs. gray), the results of which are detailed in Table 2 (see also Figure 5). In every region, there was a trend toward a crossover interaction, such that more activity was elicited in the background task by items associated with scenes than with gray backgrounds, whereas this pattern was reversed in the location task.

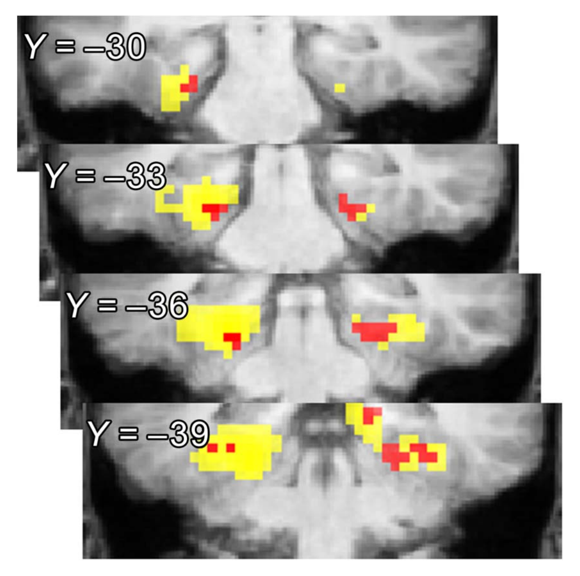

Figure 4. PHC regions identified by a voxel-wise analysis where scene reinstatement effects either did not significantly differ according to task (yellow) or demonstrated a Task × Background type interaction. See text for details.

DISCUSSION

The primary aim of this experiment was to assess whether the contents of recollection are modulated by the demands of the retrieval task. We addressed this question by exploiting the phenomenon of cortical reinstatement, taking advantage of prior findings that demonstrated that retrieval of visual scenes is associated with enhanced activity in the same cortical regions that are selectively engaged during online scene processing (Gordon et al., 2014; Johnson et al., 2013; Kuhl et al., 2013; Staresina et al., 2012; Johnson & Rugg, 2007). We assessed whether "scene reinstatement effects" varied according to whether scene information was relevant (the background task) or irrelevant (the location task) to the retrieval goal. As expected on the basis of prior findings (see above), scene reinstatement effects were evident in PHC and RSC. There was no evidence that the effects in RSC were task sensitive, suggesting that scene information was reinstated in this region to a comparable extent regardless of whether or not it was relevant to the retrieval goal. In PHC, however,

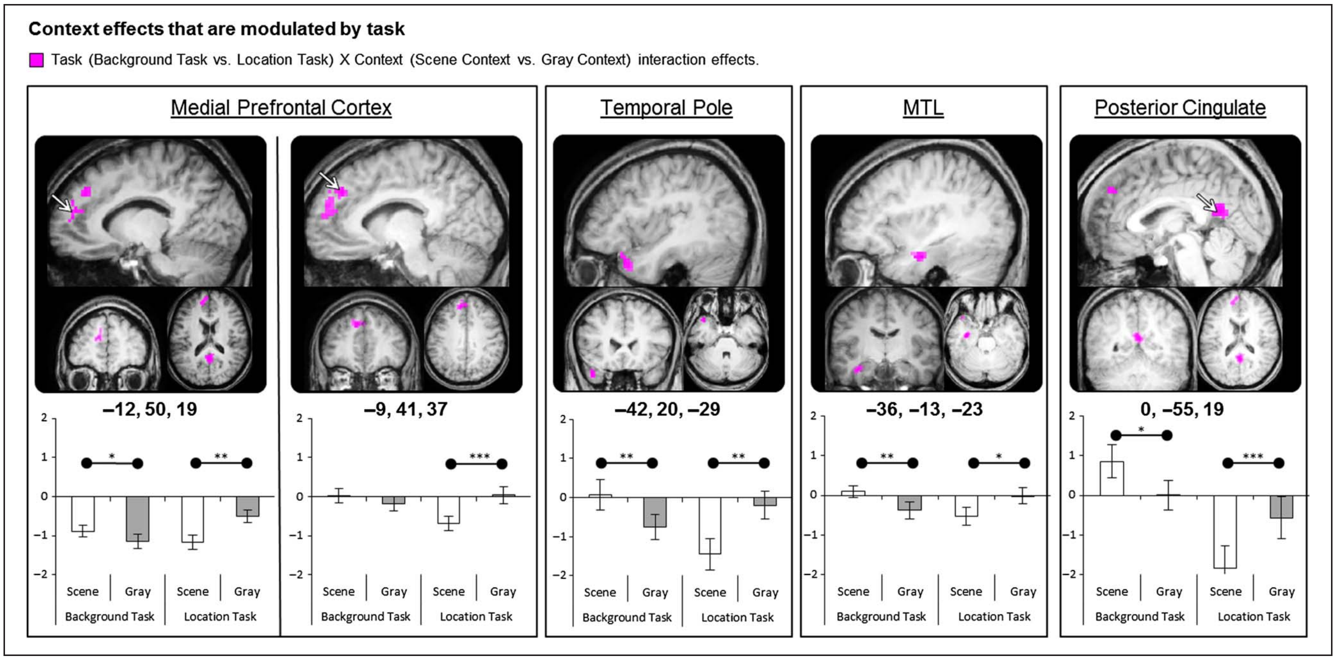

Figure 5. Whole-brain Task × Background type interaction effects. Mean parameter estimates (arbitrary units), derived from 3-mm spheres centered on the peak of each interaction effect, are illustrated below each section. Error bars indicate ±1 *SEM*. *$p < .05$, **$p < .01$, ***$p < .001$.

scene reinstatement effects were weaker when scene information was task irrelevant. Below, we discuss the implications of these findings for an understanding of how the contents of recollection are controlled to align retrieval with behavioral goals.

Behavioral Findings

We first discuss the behavioral findings and their implications in respect of the aims of the experiment. The first noteworthy aspect is that the two source memory tasks differed in their difficulty, with performance on the background task markedly exceeding that for the location task. This suggests that there was likely an asymmetry in the potential for interference from the retrieval of task-irrelevant contextual features in the two tasks. Whereas the relative inaccessibility of location information meant that its incidental retrieval would have occurred on only a minority of the trials where background information was the relevant feature, the memorability of the backgrounds suggests that this information would have been accessible, if not actually retrieved, on the majority of the location trials. Thus, to the extent that participants adopted a strategy of "suppressing" retrieval of task-irrelevant features (see below), the strategy would likely have been more useful in the location than in the background task.

The second relevant finding concerns the accuracies and RTs associated with the location judgments. Accuracy was lower, and RTs longer, for judgments performed on items studied against scene than against gray backgrounds. It is tempting to interpret this finding as evidence for retrieval interference. By this argument, participants were unable to fully prevent retrieval of a test item's studied background when cued to perform the location task.

Table 1. Regions Where Activity Associated with Source Correct Responses Shows in Interaction with Context (Scene vs. Gray) and Task (Background Task vs. Location Task)

Region	*x, y, z*	*Peak Z*	*Cluster Extent*
L. mPFC	−12, 50, 19	3.71	55 voxels
L. mPFC	−9, 41, 37	4.14	43 voxels
PC	0, 55, 19	4.11	65 voxels
L. medial-temporal lobe	−36, −13, −23	4.45	31 voxels
L. temporal pole	−42, 20, −29	3.97	28 voxels

Table 2. Simple Effects Test to Compare Activity Associated Scenes and Gray Items in Each Task

	Background Task		*Location Task*	
Region	*T*	*p*	*T*	*p*
L. mPFC	2.16	<0.05	−3.47	<0.005
L. mPFC	1.47	*ns*	−4.28	<0.001
PC	2.31	<0.05	−3.84	<0.001
L. medial-temporal lobe	2.98	<0.01	−2.46	<0.05
L. temporal pole	3.18	<0.01	−3.08	<0.01

When the retrieved background comprised scene information, it acted as a distractor, interfering with retrieval of location information to a greater extent than when the background was gray. This interpretation fits well with the proposal, outlined below, that participants actively attempted to suppress scene retrieval on location trials. The present data do not, however, allow us to rule out an alternative account of these findings, namely, that they reflect interference at encoding. By this argument, the additional attentional resources attracted by scenes relative to the gray backgrounds detracted from the resources available to encode the spatial location of the words. Prior evidence suggesting that a study item's location is encoded relatively automatically (e.g., Delogu, Nijboer, & Postma, 2012; Ellis, 1990) leads us, however, to favor the aforementioned retrieval interference account of the present findings.

Finally, a noteworthy, albeit puzzling, finding was that item memory was more accurate in the background task than in the location task. This cannot be attributed to differential encoding, because the study task preceding the two retrieval tests was identical. Presumably, the finding is a reflection of the different retrieval demands of the two tasks, but why these different demands should have impacted recognition memory is unclear.

fMRI Findings

As already noted, the primary aim of this study was to examine whether recollected content varies with the goal of the retrieval attempt. The findings suggest that, at least under the present experimental circumstances, retrieval goal does modulate the contents of recollection, but only partially. Whereas scene reinstatement effects (see below for further discussion of the interpretation of these effects) were of equivalent magnitude in the two tasks in RSC, the effects in PHC were weaker in the location task and, indeed, were not detectable in right PHC in that task.

How can this regional dissociation in the goal-dependent modulation of scene reinstatement effects be explained? An important clue comes from proposals that scene information (and, perhaps, other kinds of contextual information) are represented at different levels of abstraction in RSC and PHC. It has been proposed that information is represented at a "gist-like," relatively abstract level in RSC, but at a more fine-grained and specific level in PHC (Aminoff, Kveraga, & Bar, 2013; Bar, 2004). In light of these proposals, we suggest that the present findings are consistent with the notion that, regardless of the retrieval task, a coarse-grained representation of a studied item's background was reinstated whenever recollection occurred. When the retrieval goal depended upon more detailed information about the background (whether the information represented an urban or rural setting), retrieval of this coarse-grained representation was accompanied by retrieval of additional fine-grained information supported by PHC. By contrast, when fine-grained information about the background was goal-irrelevant and, according to one interpretation of the behavioral findings (see above), detrimental to the retrieval goal, its retrieval was gated, with a corresponding down-regulation of PHC. By this argument, therefore, episodic retrieval can operate iteratively, such that an entire study episode is not necessarily reinstated at maximum fidelity. Instead, the contents of retrieval are, metaphorically speaking, "filtered" by the retrieval goal.

Importantly, this dissociation between the RSC and PHC is incompatible with the notion that participants were capable of adopting a retrieval set or "orientation" (Rugg & Wilding, 2000) in the location task that prevented retrieval of irrelevant scene information, but that this "preretrieval" strategy succeeded on only some trials. By this account, the attenuated scene reinstatement effects in the location task were a consequence of the mixing of trials where scene recollection was successfully prevented and other trials where recollection occurred to the same extent as in the background task. Were this account correct, attenuated scene reinstatement would have been evident not only in PHC but in RSC also. As already noted, there was, however, no evidence of attenuated reinstatement in the latter region (see Results and Figure 3).

The present findings seemingly differ from those reported by Kuhl et al. (2013). As was described in the Introduction, these authors also addressed the question of whether recollection of episodic content (in their case, scenes and faces) differs depending on the task relevance of the content. They reported that, as assessed by the output of an MVPA classifier, scene and face reinstatement in PHC and fusiform cortex did not differ according to the retrieval task. Accordingly, Kuhl et al. (2013) concluded that recollected content was not modulated by retrieval goal. There are several possible reasons for the divergence between those results and the present findings. For example, as was mentioned in the Introduction, the experimental designs differ in potentially important ways (e.g., the employment of repeated study–test cycles as opposed to a single cycle). Additionally, the fMRI analysis approaches are also divergent (MVPA vs. univariate analysis). Whereas the results from these two approaches can be congruent, there is no necessity that this is the case (Davis & Poldrack, 2013). It is conceivable, for example, that intervoxel patterning of scene-related activity in PHC in the present experiment was less affected by the manipulation of retrieval goal than was mean signal and that MVPA would have not detected goal-dependent reinstatement effects. Were this to be the case, however, it would not detract from the conclusion that retrieval-related activity in this region (as indexed by mean BOLD signal) was modulated by retrieval goal.

Which brain regions might have played a role in modulating the retrieval of scene information according to its task relevance? Some clues are provided by the interaction analyses that examined where background effects (scene vs. gray) varied according to task. A reliable interaction effect was identified in four regions (see Figure 5). With the exception of one of the two mPFC clusters, where effects were confined to the location task, the interaction took the form of a crossover, with lower activity associated with scenes than the gray background in the location task and an opposite effect in the background task. As discussed below, the finding that activity elicited by items studied with scenes was attenuated in the location task may reflect the direction of attention away from scene information when it was not relevant to the retrieval task.

Importantly, mPFC and PC have both previously been implicated in a control process responsible for "suppressing" activity in PHC when scene information is task irrelevant (Chadick, Zanto, & Gazzaley, 2014; Chadick & Gazzaley, 2011). In these studies, participants saw overlapping scenes and faces under instructions to attend to and remember only one of the two stimulus categories. Relative to a passive viewing condition, activity in the PHC was reduced when scenes were task irrelevant and enhanced when they were relevant. This "suppression" of scene-related PHC activity was accompanied both by enhanced connectivity between PHC and mPFC and PC and by a reduction in mean signal in mPFC and PC. Chadick et al. (2014) proposed that these findings reflected the role of the mPFC and PC in a "suppression network" that, when down-regulated, dampened activity in cortical regions responsible for representing goal-irrelevant perceptual information. Consistent with this proposal, they reported that, in a sample of older participants, there was a negative correlation between amount of mPFC "deactivation" and the detrimental effects on face memory caused by the concurrently presented irrelevant scenes.

The proposal of Chadick and Gazzaley (2011) that reduced activity in mPFC and PC is associated with suppression of PHC lends itself well to an account of the task by background context interactions found in mPFC and PC in the present experiment. As already discussed, we suggest that to facilitate retrieval of location information, scene reinstatement was (partially) suppressed on location trials. Following Chadick and Gazzaley, we further suggest that the mechanism by which suppression was effected involved down-regulation of activity in mPFC and PC, as evidenced by the lower activity in these regions on location trials when the test item had been studied with a scene rather than the gray background (Figure 4). We speculate that the tendency for the same regions to demonstrate the opposite effect in the background task reflects a mechanism for enhancing the representation of scene information when the information was task relevant. By the account offered here, therefore, mPFC and PC played key roles in regulating the retrieval and representation of encoded scene information according to its task relevance. It remains to be determined whether these roles extend to other kinds of episodic information.

Importantly, the suppression-related activity in the mPFC and PC must reflect processes that operate post-retrieval, because there is no other basis for these context-dependent effects. This does not mean that the effects depended upon retrieval of a complete representation of the study episode but merely one of sufficient fidelity to permit discrimination between scene and gray backgrounds, such as would arguably be afforded by the RSC (Bar, 2004). We conjecture that the mPFC and PC support postretrieval (or more accurately, perhaps, "intraretrieval") processes that acted in this study to facilitate or suppress the retrieval of fine-grained scene information supported by PHC according to its task relevance.

Finally, we note that throughout the discussion above we have referred to the differences in activity in PHC and RSC elicited by items paired with scene and gray backgrounds as scene reinstatement effects. Our use of this term is supported by the evidence that these regions are activated to a greater extent by scenes than by other classes of visual input, such as faces or objects (e.g., Park & Chun, 2009; Epstein, 2008), and that prior retrieval studies have reported similar effects. It has been proposed, however, that RSC and, in particular, PHC, support the representation of not only visuospatial information but also contextual information more generally (Aminoff et al., 2013; Diana, Yonelinas, & Ranganath, 2012; Eichenbaum, Yonelinas, & Ranganath, 2007; but see Mullally & Maguire, 2011). If this proposal is correct, it might seem surprising that RSC and PHC would demonstrate differential activity according to whether a test item had been paired with a scene or a gray background: The latter backgrounds were arguably no less of a context than the scenes and, presumably, were as likely to be successfully retrieved and "reinstated." Although we would argue that the present findings are more compatible with a role for RSC and PHC in the processing of visuospatial rather than more generic contextual features, the findings are not incompatible with the alternate possibility. Unlike the scenes, which were trial-unique, the gray background did not differ across trials. If, as is the case for direct perceptual information, the RSC and PHC demonstrate "repetition suppression" when the same contextual information is repeatedly retrieved (as was reported by Diana et al., 2012), the present findings can be easily accounted for even if RSC and PHC represent both scene-based and other forms of context. Importantly, this alternate interpretation of the present findings does not require modification of our proposal that scene reinstatement in PHC varies according to the retrieval goal.

To conclude, we found evidence for scene reinstatement effects in both RSC and PHC in a retrieval task where memory for the scenes was not only task irrelevant but also arguably detrimental to task performance. These effects were, however, attenuated in PHC relative to when scenes were task relevant and the attenuated effects were accompanied by reduced activity in mPFC and PC. On the basis of these findings, we suggest that the contents of recollection are at least partially sensitive to the goal of a retrieval attempt and that representations of task-irrelevant features of an episode can be actively suppressed.

Acknowledgments

This research was supported by NIMH Grant 5R01MH072966.

Reprint requests should be sent to Rachael L. Elward, Center for Vital Longevity, 1600 Viceroy Drive, Suite 800, Dallas, TX 75235, or via e-mail: rachael.elward@utdallas.edu.

REFERENCES

Aminoff, E. M., Kveraga, K., & Bar, M. (2013). The role of the parahippocampal cortex in cognition. *Trends in Cognitive Sciences, 17,* 379–390.

Anderson, M. C., & Bjork, R. A. (1994). Mechanisms of inhibition in long term memory. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 265–325). San Diego, CA: Academic Press.

Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience, 5,* 617–629.

Chadick, J. Z., & Gazzaley, A. (2011). Differential coupling of visual cortex with default or frontal-parietal network based on goals. *Nature Neuroscience, 14,* 830–832.

Chadick, J. Z., Zanto, T. P., & Gazzaley, A. (2014). Structural and functional differences in medial prefrontal cortex underlie distractibility and suppression deficits in ageing. *Nature Communications, 5,* 4223.

Cocosco, A., Kollokian, V., Kwan, R. S., & Evans, A. C. (1997). Proceedings of the 3rd International Conference on Functional Mapping of the Human Brain. Copenhagen; Denmark: Brainweb: Online interface to a 3D MRI simulated brain database.

Danker, J. F., & Anderson, J. R. (2011). The ghosts of brain states past: Remembering reactivates the brain regions engaged during encoding. *Psychological Bulletin, 136,* 87–102.

Davis, T., & Poldrack, R. A. (2013). Measuring neural representations with fMRI: Practices and pitfalls. *Annals of the New York Academy of Sciences, 1296,* 108–134.

Delogu, F., Nijboer, T. C. W., & Postma, A. (2012). Encoding location and serial order in auditory working memory: Evidence for separable processes. *Cognitive Processing, 13,* 267–276.

Diana, R., Yonelinas, A. P., & Ranganath, C. (2012). Adaptation to cognitive context and item information in the medial temporal lobes. *Neuropsychologia, 50,* 3062–3069.

Eichenbaum, H., Yonelinas, A. R., & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annual Review of Neuroscience, 30,* 123–152.

Ellis, N. R. (1990). Is memory for spatial location automatically encoded? *Memory & Cognition, 18,* 584–592.

Epstein, R. A. (2008). Parahippocampal and retrosplenial contributions to human spatial navigation. *Trends in Cognitive Sciences, 12,* 388–396.

Friston, K. J., Holmes, P., Worsley, K. J., Poline, J.-P., Frith, C. D., & Frackowiak, R. S. J. (1995). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping, 2,* 189–210.

Friston, K. J., Penny, W., Phillips, C., Kiebel, S., Hinton, G., & Ashburner, J. (2002). Classical and Bayesian inference in neuroimaging: Theory. *Neuroimage, 16,* 465–483.

Gordon, A. M., Rissman, J., Kiani, R., & Wagner, A. D. (2014). Cortical reinstatement mediates the relationship between content-specific encoding activity and subsequent recollection decisions. *Cerebral Cortex, 24,* 3350–3364.

Halamish, V., Goldsmith, M., & Jacoby, L. L. (2012). Source-constrained recall: Front-end and back-end control of retrieval quality. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8,* 1–15.

Herron, J. E., & Wilding, E. L. (2006). Neural correlates of control processes engaged before and during recovery of information from episodic memory. *Neuroimage, 30,* 634–644.

Jacoby, L. L., Shimizu, Y., Daniels, K., & Rhodes, M. G. (2005). Modes of cognitive control in recognition and source memory: Depth of retrieval. *Psychonomic Bulletin & Review, 12,* 852–857.

Johnson, J. D., & Rugg, M. D. (2007). Recollection and the reinstatement of encoding-related cortical activity. *Cerebral Cortex, 17,* 2507–2515.

Johnson, J. D., Suzuki, M., & Rugg, M. D. (2013). Recollection, familiarity, and content-sensitivity in lateral parietal cortex: A high-resolution fMRI study. *Frontiers in Human Neuroscience, 7,* 219.

Kuhl, B., Johnson, M. K., & Chun, M. M. (2013). Dissociable neural mechanisms for goal-directed versus incidental memory reactivation. *The Journal of Neuroscience, 33,* 16099–16109.

Kuhl, B., Rissman, J., & Wagner, A. D. (2012). Multi-voxel patterns of visual category representation during episodic encoding are predictive of subsequent memory. *Neuropsychologia, 50,* 458–469.

Mullally, S. L., & Maguire, E. (2011). A new role for the parahippocampal cortex in representing space. *Journal of Neuroscience, 31,* 7441–7449.

Norman, K. A., & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. *Psychological Review, 110,* 611–646.

Park, S., & Chun, M. M. (2009). Different roles of the parahippocampal place area (PPA) and retrosplenial cortex (RSC) in panoramic scene perception. *Neuroimage, 47,* 1747–1756.

Rugg, M. D. (2004). Retrieval processing in human memory: Electrophysiological and fMRI evidence. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (3rd ed., pp. 727–739). Cambridge, MA: MIT Press.

Rugg, M. D., Johnson, J. D., & Uncapher, M. R. (2015). Encoding and retrieval in episodic memory: Insights from fMRI. In A. Duarte, D. R. Barense, & M. Addis (Eds.), *Handbook on the cognitive neuroscience of memory* (pp. 84–107). Wiley-Blackwell.

Rugg, M. D., & Wilding, E. L. (2000). Retrieval processing and episodic memory. *Trends in Cognitive Sciences, 4,* 108–115.

Snodgrass, J. G., & Corwin, J. C. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General, 117,* 34–50.

Staresina, B. P., Henson, R. N., Kriegeskorte, N., & Alink, A. (2012). Episodic reinstatement in the medial temporal lobe. *Journal of Neuroscience, 32,* 18150–18156.

Wimber, M., Alink, A., Charest, I., Kriegeskorte, N., & Anderson, M. C. (2015). Retrieval induces adaptive forgetting of competing memories via cortical pattern suppression. *Nature Neuroscience, 18,* 582–589.