

Recollection, Familiarity, and Cortical Reinstatement: A Multivoxel Pattern Analysis

Jeffrey D. Johnson,^{1,4,*} Susan G.R. McDuff,^{2,4} Michael D. Rugg,¹ and Kenneth A. Norman^{2,3}

¹Center for the Neurobiology of Learning and Memory and Department of Neurobiology and Behavior, University of California, Irvine, Irvine, CA 92697, USA

²Department of Psychology

³Princeton Neuroscience Institute

Princeton University, Princeton, NJ 08544, USA

⁴These authors contributed equally to this work.

*Correspondence: jeff.johnson@uci.edu

DOI 10.1016/j.neuron.2009.08.011

SUMMARY

Episodic memory retrieval is thought to involve reinstatement of the neurocognitive processes engaged when an episode was encoded. Prior fMRI studies and computational models have suggested that reinstatement is limited to instances in which specific episodic details are recollected. We used multivoxel pattern-classification analyses of fMRI data to investigate how reinstatement is associated with different memory judgments, particularly those accompanied by recollection versus a feeling of familiarity (when recollection is absent). Classifiers were trained to distinguish between brain activity patterns associated with different encoding tasks and were subsequently applied to recognition-related fMRI data to determine the degree to which patterns were reinstated. Reinstatement was evident during both recollection- and familiarity-based judgments, providing clear evidence that reinstatement is not sufficient for eliciting a recollective experience. The findings are interpreted as support for a continuous, recollection-related neural signal that has been central to recent debate over the nature of recognition memory processes.

INTRODUCTION

Findings from psychological and neurobiological studies of memory have led to general agreement that many of the neurocognitive processes engaged when an event is encoded are re-engaged when the event is retrieved (Damasio, 1989; Rugg et al., 2008). This *reinstatement* of encoding-related processing during retrieval is a major component of several neurally inspired models of episodic memory (e.g., Alvarez and Squire, 1994; McClelland et al., 1995; Norman and O'Reilly, 2003; Rolls, 2000; Shastri, 2002). According to the model of Norman and O'Reilly (2003), for example, the neural architecture of the hippocampus allows it to store nonoverlapping representations of the

patterns of cortical activity elicited when different events are encoded. When an effective retrieval cue for an event is presented, the appropriate hippocampal representation is reactivated, leading to reinstatement of the original pattern of cortical activity. Crucially, in the context of “dual-process” theories of recognition memory (Mandler, 1980; for review, see Yonelinas, 2002), hippocampally mediated reinstatement is thought to support the *recollection* (or recall) of specific details associated with an episode. In contrast to the hippocampal memory system, a cortical system involving extrahippocampal regions of the medial temporal lobe is capable of giving rise only to an acontextual (nonrecollective) *familiarity* signal that corresponds to the scalar match between the cue and episode (Norman and O'Reilly, 2003).

With the exception of indirect neuropsychological evidence (e.g., Rubin and Greenberg, 1998), empirical support for recollection-related cortical reinstatement in humans comes largely from the use of functional neuroimaging. Studies employing event-related functional magnetic resonance imaging (fMRI) have been particularly useful in this regard by enabling the neural correlates of recollection to be contrasted according to the nature of the recollected content (e.g., Johnson and Rugg, 2007; Kahn et al., 2004; Wheeler and Buckner, 2004; Woodruff et al., 2005). In combination with the neural measure, these studies adopted behavioral procedures—such as the *remember/know* (Tulving, 1985) or source memory procedures (Johnson et al., 1993)—to identify trials where subjects retrieved specific episodic details. (To minimize confusion, we hereafter use the term *remembering* to refer to the experience of retrieving specific episodic details, and we use the term *knowing* to refer to the experience of recognizing an item without retrieving specific episodic details. We reserve the terms *recollection* and *familiarity* for describing the processes and neural signals often thought to respectively underlie those experiences.) Arguably the most convincing evidence from these studies in favor of reinstatement comes in the form of regionally specific double dissociations in the cortical patterns associated with the remembering differential content. For instance, Kahn et al. (2004) reported that remembering words studied in a visual imagery task activated left parahippocampal cortex to a greater extent than did remembering words studied in a phonological task, whereas the reverse contrast was associated with activation of left premotor cortex.

Similarly, Woodruff et al. (2005) reported that two regions of left fusiform cortex—shown previously to be functionally specialized for the processing of visually-presented words and pictures—exhibited dissociable activity with respect to remembering words versus pictures.

In a direct test of the reinstatement hypothesis, Johnson and Rugg (2007) investigated the extent to which content-specific neural correlates of remembering overlapped with regions that were selectively active when the relevant content was encoded. Subjects were first presented with a series of words and required to use the words either in a sentence or in a visual imagery task and then undertook a remember/know test. Brain regions where greater activity was associated with remember compared to know responses exhibited specificity according to the class of study episode: words studied with the sentence task elicited greater activity in medial prefrontal cortex, whereas words studied with the imagery task elicited greater activity in occipital and fusiform cortex. Importantly, the regions demonstrating these dissociations were a subset of regions exhibiting differential activity when the two classes of words were initially studied (see Kahn et al., 2004, for a similar, across-experiment comparison). Thus, these findings established a direct link between the neural correlates of study processing and the phenomenological experience of remembering, consistent with the notion that processing selectively engaged during encoding is reinstated during retrieval.

Although the aforementioned findings convincingly demonstrate a relationship between cortical reinstatement and recollection, they do not address two important questions about the status of reinstatement effects when items are reportedly judged on the basis of familiarity (known). First, are such judgments at all associated with reinstatement? In the three fMRI studies described above, although behavioral methods designed to separate recollection and familiarity were employed, in two of the studies there were insufficient numbers of trials to evaluate reinstatement when know responses were given (Johnson and Rugg, 2007; Woodruff et al., 2005), while the relevant contrasts were not reported in the remaining study (i.e., for incorrect source judgments; Kahn et al., 2004). If it transpires that reinstatement effects are evident during know responses, a second question arises: do these effects differ from those associated with remembering, either in magnitude or localization?

Resolution of these two questions has important implications for the ongoing theoretical debate about the nature of processes contributing to recognition memory. According to some dual-process theorists (see Yonelinas, 2002), remember and know responses reflect the influence of qualitatively distinct processes. Remember responses are thought to reflect recollection of specific details, whereas know responses are thought to be based on familiarity (in the absence of recollection). Accepting this assertion at face value leads to the prediction that cortical reinstatement will be present for remember responses but absent for know responses. Contrary to this view, another class of theories posits that different recognition judgments are not based on a clear-cut distinction between two memory processes or signals. Rather, the judgments result from assessing a single, continuous “memory strength” signal (e.g., Donaldson, 1996; Dunn, 2004) or an amalgam of continuous signals (Wixted and Stretch, 2004;

Wixted, 2007) present across all of the different judgments. That is, each test item is associated with a particular level of the strength signal, and a detection process is used to decide whether the strength exceeds a criterion, thus determining the response. If it is assumed that the degree of retrieval-related reinstatement co-varies with memory strength, then reinstatement effects should follow a graded profile: largest when subjects report remembering episodic details yet also present for know responses that the subject attributes to familiarity.

The present fMRI study was designed to explore the relationship between cortical reinstatement and distinct phenomenological bases of recognition memory, as evidenced by different behavioral correlates (i.e., remember versus know). Subjects first completed a study phase where they viewed a series of words and undertook three different encoding tasks that elicited distinct patterns of cortical activity (*Artist*, *Function*, and *Read*; McDuff et al., 2009). During a later test phase, recognition memory for the studied words was assessed using a modified remember/know procedure, in which one of five responses was required to each test item (Yonelinas et al., 2005). One response was used to indicate that details associated with studying an item were remembered. The remaining four responses were used to rate the confidence with which an item was known to be studied or not studied, presumably on the basis of item familiarity in the absence of recollection. fMRI data acquired during both the study and test phases allowed for direct comparison between encoding- and retrieval-related activity (Johnson and Rugg, 2007). According to the view that remember (but not know) judgments veridically index retrieval of specific episodic details, reinstatement effects should be confined to items endorsed as remembered. Alternatively, according to theories positing a memory strength continuum, reinstatement effects should be evident, albeit in weaker form, for test items associated with know judgments.

In contrast to the previous studies of reinstatement described above, we employed multivoxel pattern analyses (MVPA) of the fMRI data (for reviews, see Haynes and Rees, 2006; Norman et al., 2006). MVPA is well suited for characterizing reinstatement because it quantifies the relationship between patterns of brain activity acquired during one experimental phase (the study phase in our case) and any “reactivated” patterns from another phase (our test phase). Moreover, because MVPA involves classifying correlated patterns of activity across multiple voxels, it is often considered to be more sensitive than “mass-univariate” fMRI analyses, which might fail to detect differences in signals that are weak at the single-voxel level or even when spatially smoothed across voxels (see Haynes and Rees, 2006; Norman et al., 2006). In the present study, two types of MVPA were implemented. The first type was designed to maximize the sensitivity of detecting reinstatement across different recognition memory judgments, by making use of a subset of voxels that best distinguished between the encoding tasks (for similar implementations, see McDuff et al., 2009; Polyn et al., 2005). The second type of MVPA involved classifying data from “searchlights” (spheres) of voxels (Kriegeskorte et al., 2006; Mur et al., 2009) and provided information about whether the spatial distribution of reinstatement effects throughout the brain differed according to the type of memory judgment.

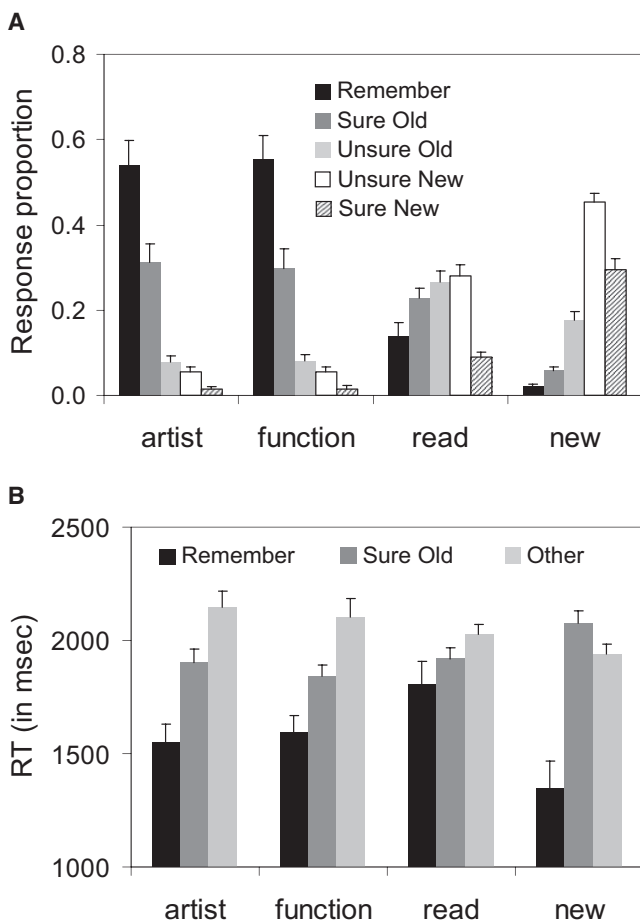


Figure 1. Behavioral Performance

(A) Mean (+SEM) proportions of responses according to the test item condition.

(B) Mean (+SEM) response time (RT) data. The Other category reflects collapsed Unsure Old, Unsure New, and Sure New responses (due to low individual trial numbers). The RT data for Remember responses to new items are based on only 12 subjects contributing such responses.

RESULTS

Behavioral Performance

Figure 1 displays the mean proportions of responses (*Remember*, *Sure Old*, *Unsure Old*, *Unsure New*, and *Sure New*) and corresponding RTs to items presented during the test phase. As is apparent in the figure, items previously studied with the Artist and Function tasks primarily elicited Remember responses, followed by Sure Old (know) responses; items studied with the Read task were associated mostly with Sure Old, Unsure Old, and Unsure New responses; and new items primarily elicited New responses. Because the pattern-classification analyses (see below) are restricted to old items, we also focused the behavioral analyses on those items. Additionally, given the low proportions of Unsure Old, Unsure New, and Sure New responses for Artist and Function items, the corresponding trials were collapsed into an *Other* category for each task. ANOVA of the response proportions, incorporating factors of task (Artist,

Function, and Read) and response (Remember, Sure Old, and Other), revealed a significant interaction ($F_{2,3,34.3} = 64.05$, $p < 0.001$; degrees of freedom corrected according to Greenhouse and Geisser, 1959). The interaction indicated a tradeoff between response categories, such that Artist and Function items elicited more Remember responses than did Read items (min. $t_{15} = 7.73$, $p < 0.001$), whereas the opposite was true for Other responses (min. $t_{15} = 13.15$, $p < 0.001$). ANOVA of the RT data (including the same factors as above) gave rise to a significant main effect of response ($F_{1,4,21.7} = 26.51$, $p < 0.001$) and its interaction with task ($F_{3,2,47.5} = 7.38$, $p < 0.001$). Subsidiary ANOVAs revealed a task effect only for Remember responses ($F_{1,4,20.4} = 12.25$, $p < 0.005$), whereby RTs were shorter for Artist and Function items than for Read items (min. $t_{15} = 3.45$, $p < 0.005$). Notably, the response proportions and RTs for Artist and Function items were statistically equivalent.

Whole-Brain MVPA

As described in the Introduction, our primary aim was to determine the relationship between patterns of brain activity elicited at study and those elicited at test. MVPA was employed to provide a sensitive index of the strength of the study-test relationship, but this index is blind to the loci of voxels expressing any such relationship. Although the specific brain regions discriminating between the three study tasks (as determined through classifier training) were largely inconsequential, it was important to ensure that the voxels were biologically meaningful—that is, they constituted sizable clusters rather than dispersed individual voxels and encompassed cortical areas expected to be active in cognitive tasks such as those employed here. Accordingly, we created a group mean *importance map* for each study task condition, by combining the voxelwise input values and the trained classifier's weights (see *Experimental Procedures*; McDuff et al., 2009). As shown in Figure 2, it is apparent that the patterns of important voxels—clustering largely in bilateral occipital, superior parietal, and left inferior frontal cortex—meet our aforementioned criteria for meaningfulness.

The accuracy of the classifier in determining the prior encoding history of old test items was operationalized as the probability that the classifier's output for the correct study task was greater than the output for each of the other two tasks. Classifier accuracy for a given test item was assessed beginning with the time point (TR) in which the test item onset (hereafter, TR 1) and continuing for six additional time points (up to TR 7). Thus the accuracy measure provided information about classification performance as a function of time. Overall classifier accuracy for all of the old items, regardless of the response given, is shown in Figure 3. As is clear from the figure, classifier accuracy was maximal at TR 4, coinciding with the expected peak of the hemodynamic response for a transient stimulus (as estimated by convolving hemodynamic and impulse response functions). To test whether classifier accuracy exceeded chance (0.33) at any of the TRs, a series of one-sample t tests was conducted. Accuracy was significantly above chance for TRs 3 through 7 (min. $t_{15} = 3.75$, $p < 0.005$), and each of these effects remained significant following correction for multiple comparisons across the seven TRs (using the Holm-Bonferroni procedure with an overall




Figure 2. Importance Maps

Group mean importance maps for the three study tasks, overlaid on axial slices of the mean normalized anatomical data (coordinates in Talairach space). The colored areas depict voxels where importance values exceeded arbitrary thresholds of 0.001 positively (red) and -0.001 negatively (green; see middle row, right-most column). L = left.

$p < 0.05$; Holm, 1979). These findings demonstrate that the patterns of fMRI data associated with old test items differed systematically, so as to allow the classifier to accurately determine an item's prior encoding history.

Of more relevance to our primary aim was whether the classifier's ability to correctly identify the encoding condition of test items was limited to items eliciting a recollection response or whether above-chance accuracy also extended to items accorded nonrecollective responses. To address this issue, we investigated classifier accuracy according to response category (Remember, Sure Old, and Other). Classifier accuracy was first assessed separately for the three tasks and then averaged across tasks. The resulting accuracy values for each response category are shown in Figure 3, in which it is apparent that accuracy once again followed the expected (hemodynamically corrected) time course. One-sample *t* tests (corrected for multiple comparisons as before) revealed that accuracy was above chance when test items were endorsed with either a Remember or Sure Old response. Significant effects were observed at TRs 3 through 5 for Remember responses (min. $t_{15} = 3.26$, $p < 0.01$), and at TRs 4 and 5 for Sure Old responses (min. $t_{15} = 4.30$, $p < 0.001$). There were no significant effects for Other responses. Thus, the classifier was capable of assigning test items to the appropriate encoding condition only when the items elicited a phenomenological sense of recollection (remembering) or high-confidence familiarity (knowing).

Having demonstrated that reinstatement was above chance for both Remember and Sure Old responses, we set out to investigate the relative strength of reinstatement for Remember versus Sure Old trials. To accomplish this goal, we switched from computing classifier accuracy to measuring classifier output strength: the real-valued output for the classifier node represent-

ing the actual (true) task condition. The key advantage of using classifier output strength is that it tracks the raw magnitude of reinstatement on each trial, whereas the accuracy measure computes a binary score for each trial (based on whether the actual task output is higher than the other task outputs) and discards information about the actual magnitude of reinstatement. We restricted these analyses of reinstatement strength to the data from TRs 3 through 5, based on our earlier findings that classifier accuracy was maximal during this time period.

As with our classifier accuracy analysis, the classifier output values for each response type were first averaged across trials within each study task and then averaged across tasks. This averaging procedure was especially important here because (as mentioned above) there were significant across-task differences in responding. That is, items from the Artist and Function tasks elicited more Remember responses and fewer Other responses than did items from the Read task. This discrepancy raises the possibility that effects of response type (e.g., Remember versus Other) on classifier output strength will be confounded with effects of task (Artist/Function versus Read). Our averaging procedure eliminates this potential confound by ensuring that each task is equally represented within each response type.

Figure 4A shows the output values for each response, averaged over the Artist, Function, and Read conditions. As can be seen in the figure, the output for Remember responses is higher than that for Sure Old (except at TR 5), which is in turn higher than for Other responses. The output values were subjected to pairwise comparisons between response categories. For the Remember versus Sure Old comparison, *t* tests revealed no significant differences. For the comparison of Sure Old versus Other, there was a significant difference at TR 5 ($t_{15} = 2.71$, $p < 0.025$; corrected for multiple comparisons). Finally, the Remember versus Other comparison gave rise to a significant difference at TR 4 ($t_{15} = 2.96$, $p < 0.01$). Thus, although the output values appeared to follow a graded profile across responses, the differences did not consistently reach significance.

We hypothesized that the weak results of the previous analyses might be attributable to the low number of Remember trials in the Read condition. For reasons described above, our analysis weighted Artist, Function, and Read trials equally when estimating reinstatement for each response type. However, the actual number of Read-task Remember trials was extremely

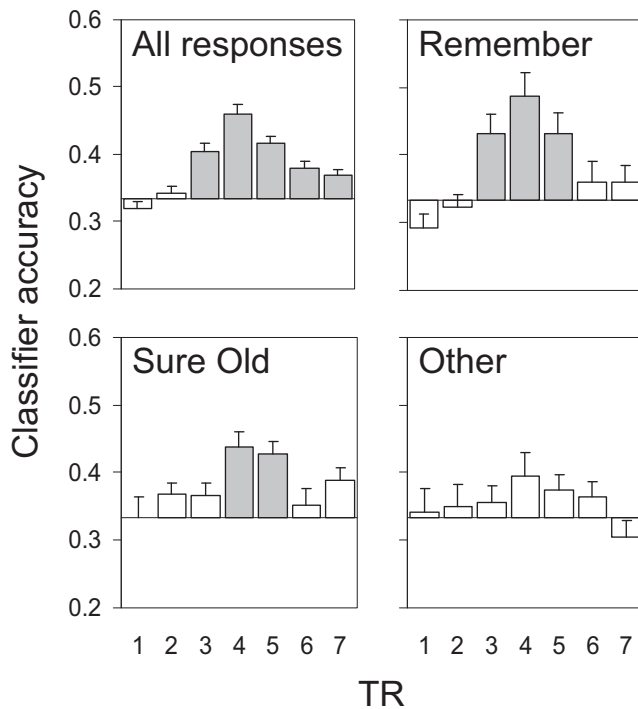


Figure 3. Classifier Accuracy

Mean classifier accuracy (+SEM) collapsed across all response categories and separated by response category. Time point (TR) 1 corresponds to test item onset. Shaded bars indicate the TRs during which classifier accuracy was significantly above chance (0.33; correcting for multiple comparisons).

small. (Each subject had at least one Remember response in the Read condition, but most subjects had fewer than five such responses.) Classifier estimates based on very small numbers of trials can be highly volatile. To address this issue, we conducted a further analysis that was identical to the foregoing one but was restricted to test items associated with the Artist and Function tasks (since these tasks, unlike the Read task, had adequate numbers of trials in each response bin). As in the previous analysis, for each response type, we first averaged classifier estimates within each task and then averaged these estimates across tasks. This procedure ensures that the Artist and Function tasks are equally represented within each response type, so effects of response type on reinstatement can not be attributed to task differences. (We also reanalyzed the accuracy data based on only the Artist and Function conditions, which produced qualitatively similar results to those reported here [see [Supplemental Data](#) available online].)

The output values averaged over the Artist and Function conditions are shown in [Figure 4B](#), segregated according to response. For the Remember versus Sure Old comparison, pairwise *t* tests revealed a significant difference for each of the three TRs (min. $t_{15} = 2.13$, $p < 0.05$; corrected for multiple comparisons). The Remember versus Other comparison also gave rise to significant differences for all TRs (min. $t_{15} = 2.85$, $p < 0.05$). In addition, the comparison of Sure Old versus Other responses revealed a significant effect at TR 5 ($t_{15} = 3.18$, $p < 0.001$). In contrast to the results based on all three study tasks, these




Figure 4. Classifier Output

Mean values (+SEM) of the classifier's correct output node, (A) averaged over all three study tasks, and (B) over only the Artist and Function tasks. Each bar reflects classifier output for a given response category and time point (TR). Brackets indicate significant differences between responses (correcting for multiple comparisons).

results clearly demonstrate a graded profile of classifier output across response categories. Specifically, output was highest for items eliciting Remember responses, intermediate for Sure Old responses, and lowest for Other responses.

Searchlight MVPA

The foregoing results provided evidence that the magnitude of reinstatement differs across recognition memory judgments. In the [Introduction](#), a further question was posed about whether the spatial patterns of reinstatement effects associated with recollection- versus familiarity-based memory judgments also differ. This question was addressed with searchlight-based classification analyses. (For comparison, a parallel GLM-based analysis is reported in the [Supplemental Data](#).) Based on our previous results demonstrating reinstatement for both Remember and Sure Old responses, we tested several possibilities regarding how the patterns of brain regions exhibiting reinstatement effects might differ according to response. First, regions might exhibit equivalent levels of above-chance reinstatement for both response types. Second, reinstatement might differ

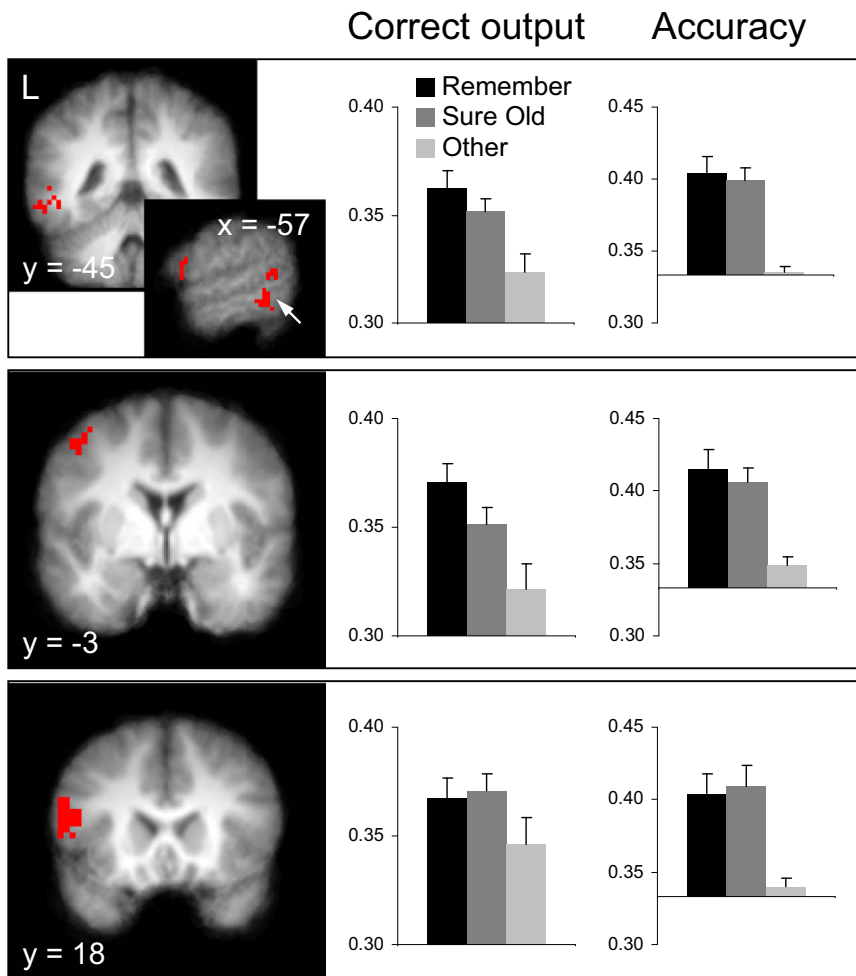


Figure 5. Equivalent Reinstatement Effects

Results of searchlight analyses where reinstatement was equivalent for test items designated with Remember and Sure Old responses (see main text for details of the contrast procedure). Histograms reflect the mean (+SEM) output values at the correct classifier node (left column) and classifier accuracy (right column; chance = 0.33) within the depicted clusters in lateral temporal cortex, superior frontal gyrus, and inferior frontal gyrus. All effects depicted here survived a cluster-wise threshold of $p < 0.05$ and are overlaid on the mean anatomical image (coordinates in Talairach space). L = left.

related differences in reinstatement magnitude also showed above-chance reinstatement.

Regions exhibiting equivalent reinstatement for Remember and Sure Old responses were identified by the intersection (identified by inclusive masking) of voxels showing above-chance classifier accuracy for the two response categories: $\text{Remember}_{\text{acc}} > 0.33$ and $\text{Sure Old}_{\text{acc}} > 0.33$ (each thresholded at $p < 0.01$). Further, any voxels where the correct classifier output differed according to the Remember versus Sure Old contrast (thresholded liberally at $p < 0.1$) were removed via exclusive masking. Figure 5 shows the outcome of this contrast procedure, which identified regions of left lateral temporal cortex, superior frontal gyrus, and inferior frontal

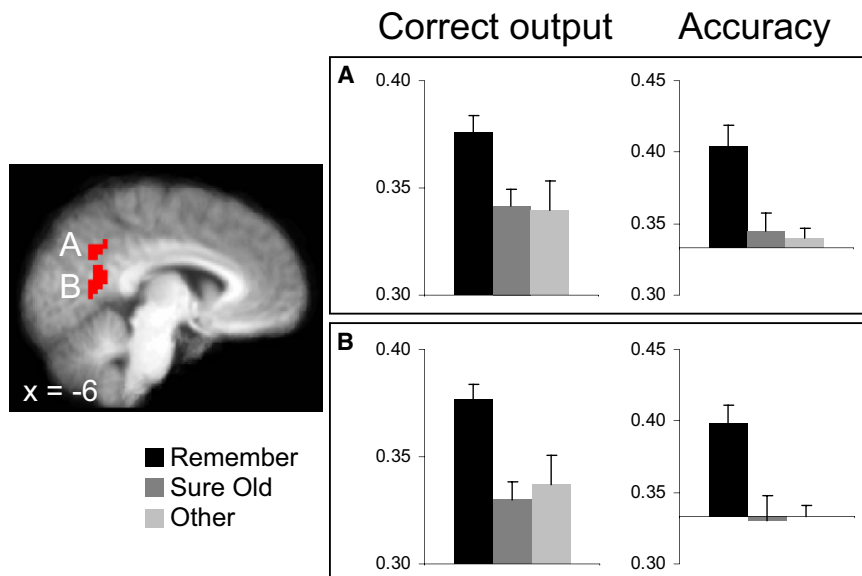
gyrus (each surviving a cluster-level correction of $p < 0.05$; [Worsley et al., 1996](#)). Thus, these regions showed reinstatement effects during both recollective- and familiarity-based memory judgments.

There were no supra-threshold clusters where reinstatement was at above-chance levels for both Remember versus Sure Old responses, but where these effects also differed in magnitude. This pattern of results was tested by contrasting the classifier output for the two responses (Remember > Sure Old or Sure Old > Remember) in combination with verifying that reinstatement was evident for both ($\text{Remember}_{\text{acc}} > 0.33$ and $\text{Sure Old}_{\text{acc}} > 0.33$; each at $p < 0.01$).

Finally, qualitatively different patterns of reinstatement for Remember and Sure Old responses were identified by testing for selective effects associated with either response. Regions exhibiting selective Remember-related reinstatement were identified with the Remember > Sure Old contrast of raw classifier output, while ensuring that reinstatement for the former response category in these voxels also achieved above-chance accuracy ($\text{Remember}_{\text{acc}} > 0.33$, each thresholded at $p < 0.01$). Additionally, any voxels where reinstatement for Sure Old responses differed from chance ($\text{Sure Old}_{\text{acc}}$ versus 0.33, bidirectional; $p < 0.1$) were excluded. As shown in Figure 6, two

quantitatively for recollection- and familiarity-based responses, whereby both are associated with above-chance reinstatement, but the effects occur at a greater magnitude for one of the responses. Finally, the patterns of reinstatement could differ qualitatively, such that the reinstatement exhibited in some regions is evident selectively for one of the responses but absent for the other response.

Given that reinstatement was previously shown to be most prominent at TRs 3 through 5, the searchlight results were simplified by averaging the classifier output values over these time points (rather than creating separate maps for each TR). For reasons outlined earlier, only the data from Artist and Function test trials were used for these analyses (first averaged separately, and then across tasks). Two types of maps—output and accuracy—were created from these results. Output maps corresponded to the real-valued output from the actual (true) task node for a given trial and were used to identify differences in the magnitude of reinstatement according to the designated test response. Accuracy maps were constructed for each response category by determining whether the output value for the actual task node was greater than that for the other two nodes (one of which was the Read node). The two map types were used together to ensure that any voxels exhibiting response-

**Figure 6. Selective Reinstatement Effects**

Results of searchlight analyses showing selective reinstatement for test items designated with Remember responses (compared to Sure Old responses; see main text for contrast procedure). The histograms provide the mean (+SEM) output value at the correct classifier node and the mean classifier accuracy within the depicted clusters of (A) posterior cingulate and (B) retrosplenial cortex. Both effects survived a cluster-wise threshold of $p < 0.05$. See Figure 5 caption for further display details.

clusters of voxels in medial posterior cortex were identified, one in the vicinity of retrosplenial cortex and the other in posterior cingulate. The analogous contrast procedure used to identify selective reinstatement for Sure Old responses (Sure Old > Remember and Sure Old_{acc} > 0.33, excluding Remember_{acc} > 0.33) revealed no significant effects. Thus, there was a single dissociation in the reinstatement effects associated with recollection- and familiarity-based memory, which took the form of regions showing selective reinstatement for recollection.

DISCUSSION

The aim of the present study was to elucidate the relationship between reinstatement of encoding-related neural activity during retrieval and different phenomenological correlates of recognition memory judgments (specifically, judgments associated with remembering versus knowing that a test item was previously encountered). Using two forms of multivariate pattern-classification analyses of fMRI data, we assessed the extent to which patterns of brain activity associated with retrieval can be used to correctly classify the prior encoding history of test items. These analyses demonstrated that MVPA is capable of detecting the relationship between brain patterns activated during encoding and those that are reactivated at test (also see McDuff et al., 2009; for analogous findings in free recall, see Polyn et al., 2005). Two novel and theoretically substantive findings emerged: one involved the different levels of reinstatement that were associated with recognition judgments having distinct subjective bases, and the other concerned the cortical regions that exhibited these reinstatement effects when the different types of judgments were made. We discuss these findings in turn below.

Using a whole-brain MVPA approach designed to be maximally sensitive to detecting reinstatement effects, we were able to classify with above-chance accuracy the prior encoding task that was undertaken for a given test item, regardless of whether the item was correctly judged as old in association

with a Remember or a Sure Old response. The results for Remember responses are consistent with findings from a recent study where it was demonstrated that the neural correlates of Remember responses overlapped with regions that were selectively active when the test items were initially studied (Johnson and Rugg, 2007; also see Kahn et al., 2004). In keeping with our prior interpretation, the present findings are taken to indicate the reinstatement of study content at the time of retrieval. The reinstated content likely reflects a recapitulation of the cognitive operations that were engaged by the different tasks during the study phase, given that there were no physical differences between test items that correlated with their prior study task.

The present findings also constitute a theoretically important extension to our prior conceptualization of reinstatement. As was noted in the Introduction, reinstatement (or content specificity) was previously evaluated for only those recognized test items that were accompanied by either a Remember response (Johnson and Rugg, 2007; Woodruff et al., 2005) or a correct source memory attribution (Kahn et al., 2004). Here, however, we have demonstrated that reinstatement is also evident for test items correctly recognized in the absence of any avowed retrieval of specific episodic details. More specifically, the magnitude of reinstatement, as measured by classifier output values, decreased in a graded manner across Remember, Sure Old, and Other responses. Therefore, although reinstatement has been shown here, as previously, to be correlated with the phenomenal experience of remembering, the current study provides compelling evidence that reinstatement is not uniquely associated with such an experience.

In a second set of analyses, we employed searchlight-based classifications to characterize the similarities and differences among reinstatement effects associated with Remember versus Sure Old responses. These analyses yielded two results. First, multiple regions exhibited reinstatement for both response types, with common reinstatement effects evident in left-lateralized regions of inferior frontal gyrus, superior frontal gyrus, and lateral temporal cortex (Figure 5). Second, an additional set of regions was associated with reinstatement for Remember responses but not for Sure Old responses. These selective effects were in the vicinity of retrosplenial cortex and posterior cingulate (Figure 6). Thus, the searchlight analyses demonstrated that the

reinstatement effects associated with Sure Old responses were in a subset of the cortical regions associated with Remember-related reinstatement.

The above findings, together with the graded reinstatement effects described earlier, suggest that the Remember and Sure Old responses relied on a common process or signal. An obvious account of the findings is that test items gave rise to different levels of a continuous memory signal, and that criteria placed along this continuum were used to assign items to the different response categories (Donaldson, 1996; Dunn, 2004). By this argument, Remember responses were made when test items evoked a signal that exceeded the strictest criterion, Sure Old responses resulted from a memory signal that fell between this criterion and one that was less strict, and Other responses were due to the signal falling short of both criteria. A similar account has been applied to results from a recent behavioral study that combined remember/know and source memory judgments (Wais et al., 2008). In that study, subjects' source memory performance was above chance even for items they reported to not remember, leading the authors to suggest that the retrieval of source information contributed to the memory strength of the resulting know responses. This is not to say that recognition memory judgments are guided solely by a single process; instead, one can assume the involvement of multiple continuous processes, whereby all types of judgments are influenced to some degree by each process (Wixted, 2007; Wixted and Stretch, 2004). Importantly, the reinstatement effects for both Remember and Sure Old responses in the present study are inconsistent with models in which such responses are thought to selectively tap into qualitatively different memory processes (i.e., recollection versus familiarity, respectively; Yonelinas, 2002). Our findings provide crucial evidence that cortical reinstatement effects constitute a neural signature of previously hypothesized instances of "subthreshold" recollection.

As described in the Introduction, our work here was largely inspired by a framework in which learning and memory rely on complementary systems: a hippocampal system capable of rapidly encoding nonoverlapping conjunctions of the cortical patterns that represent specific episodes, and an extra-hippocampal (cortical) system that exploits overlapping representations of the general statistical structure evident across similar episodes (McClelland et al., 1995; O'Reilly and Norman, 2002; O'Reilly and Rudy, 2001). Models derived from this framework, along with related models of hippocampal function, have proposed that a hippocampally stored cortical representation mediates the reinstatement of a corresponding cortical pattern, leading to recollection (Alvarez and Squire, 1994; Norman and O'Reilly, 2003; Rolls, 2000; Shastri, 2002). Consistent with this proposal, and with our findings of graded reinstatement effects, a GLM analysis (see Supplemental Data) gave rise to greater right hippocampal activity associated with Remember compared to Sure Old responses. By contrast, activity in hippocampus was not enhanced for Sure Old relative to Other responses (or correctly rejected new items). Although this latter finding might suggest that the reinstatement observed for Sure Old responses is attributable to some mechanism other than hippocampally-mediated recollection, the result is likely due instead

to the hippocampus being involved additionally in novelty related encoding processes (also see Düzel et al., 2003; Stark and Okado, 2003; Stern et al., 1996). Such processing is elicited to a greater extent by relatively unfamiliar items (e.g., new items and those given Other responses) and would thereby counteract any reinstatement-related enhancement of hippocampal activity elicited by familiar items.

As we alluded to above, a parsimonious account of the present results, in relation to the Norman and O'Reilly (2003) model, supposes that the mapping between the neural correlates of hippocampally mediated recollection and subjects' behavioral responses is more continuous than sometimes conceptualized by dual-process models of memory. It is important, however, not to overlook a discontinuity between the reinstatement effects associated with items endorsed as recollected versus those related to confident old judgments. Although reinstatement-related neural signals associated with Sure Old responding were observed in multiple cortical regions, the effects were evidently insufficient to support the phenomenal experience of recollection. At the moment, it is not possible to discern between two explanations of this discontinuity. On the one hand, as evidenced by the effects in additional cortical regions for Remember compared to Sure Old responses, recollection might result from a quantitative increase in either the number of regions exhibiting reinstatement or the magnitude of reinstatement in those areas. Alternatively, it is possible that the specific loci of reinstatement effects associated with Remember responses, such as in medial parietal cortex, carry crucial qualitative information that drives the episodic evidence above the appropriate decision threshold. Importantly, while the interpretation of this response-related distinction in the neural signal is an important topic for follow-up research, its resolution does not detract from our main finding that reinstatement plays a role in phenomenologically distinct forms of recognition memory.

Two caveats to the interpretation of the present findings deserve further discussion. First, the precise time course of the neural events driving reinstatement effects cannot be determined by fMRI data alone (for similar discussion, see Johnson et al., 2008; Johnson and Rugg, 2007; Kahn et al., 2004; Maratos et al., 2001; Woodruff et al., 2005). On the one hand, the effects might occur shortly after test item onset, as would be expected if they were a key determinant of the recognition memory judgment. Alternatively, the effects could be a consequence of the memory judgment, possibly reflecting the deployment of attention toward particular types of retrieved content, or the maintenance of that content in working memory in service of further evaluation. These two accounts can only be adjudicated by employing a neural measure with much higher temporal resolution than fMRI, such as event-related potentials (ERPs; for an example of content-specific ERP effects during retrieval, see Johnson et al., 2008). A second caveat is that the classifier was trained solely to detect patterns of activity that discriminated between the three encoding tasks. This training procedure gives the classifier the ability to detect when test items are accompanied by activity related to the recollection of task-specific details. However, the procedure does not enable the classifier to detect activity related to the recollection of "nondiagnostic" details (i.e., details shared by all three tasks) or activity associated to

familiarity-based processing—both of which likely contribute to some degree to subjects' responses in this task.

To conclude, the fMRI findings reported here are consistent with the idea that the retrieval of episodic memories involves reinstating patterns of cortical activity that were engaged during encoding. The present findings extend previous results by demonstrating that reinstatement is not restricted to instances in which subjects reportedly retrieve specific episodic information, emphasizing that the presence of a content-dependent neural signal is not sufficient for eliciting a phenomenological sense of remembering (and the ensuing response). Rather, in situations where recognition is indicated as being guided solely by a strong feeling of familiarity or knowing, reinstatement is also evident (albeit at a lower magnitude) and recruits largely the same pattern of brain regions that were associated with remember-related effects. Finally, the current study adds to a growing body of evidence demonstrating the benefits of using multivariate classification analyses to detect subtle, yet informative patterns in fMRI data.

EXPERIMENTAL PROCEDURES

Subjects

Sixteen volunteers (11 females) between 18 and 31 years of age ($M = 22$) were recruited from the undergraduate and graduate student community of Princeton University and remunerated for their participation. All subjects reported being right-handed, native-English speakers with normal or corrected-to-normal vision, no history of neurological disease, and no other contraindications for MRI. Informed consent was obtained in accordance with the Princeton University Institutional Review Board guidelines.

Stimuli

The stimuli were 306 words drawn from the MRC database (Coltheart, 1981; Wilson, 1988; http://www.psy.uwa.edu.au/mrcdatabase/uwa_mrc.htm). The words were between four and nine letters long ($M = 5.5$, $SD = 1.3$), had written frequencies between one and 50 per million ($M = 16.9$, $SD = 13.2$; Kucera and Francis, 1967) and had scores of at least 500 on scales of familiarity ($M = 580.9$, $SD = 34.6$), concreteness ($M = 539.2$, $SD = 27.5$), and imagability ($M = 581.7$, $SD = 31.3$). Words with emotional connotations or referring to alcoholic beverages were not used. Twenty-seven additional words with similar characteristics served as practice stimuli. All words were displayed visually in white lowercase 30-point Helvetica font.

For each subject, 162 words were randomly selected from the pool to serve as study items. Study items were randomly assigned to three study blocks and three encoding task conditions (see below), resulting in 18 items per task per block. A subset of 144 study items (16 drawn from each task/block combination) were re-presented as old items during the test phase and intermixed with the remaining 144 nonstudied (new) words. The test items were randomly divided into four test blocks (72 items per block), with the constraint that each block had equal numbers of old and new items.

Behavioral Tasks

Subjects were instructed and completed a practice version of the experiment outside the scanner. In the scanner, the experiment consisted of three study blocks followed by four test blocks, with each block corresponding to a separate run of fMRI acquisition. Blocks were separated by breaks of around 1–2 min. A longer (~10 min) break occurred between the study and test phases, during which anatomical data were acquired while a nature video was shown. All experimental stimuli were displayed on a screen positioned at the head of the magnet bore, which was viewed through a mirror placed in front of the subject's eyes.

For the study phase, subjects were presented with a series of words and had to complete an *Artist*, *Function*, or *Read* task for each word (Davachi et al.,

2003; Dzulkipli and Wilding, 2005; Johnson et al., 1997; see McDuff et al., 2009, for use of the same tasks). For the *Artist* task, subjects were to imagine how an *Artist* would draw the object denoted by the word and then rate the difficulty of drawing (1 = easy to 5 = hard). For the *Function* task, subjects had to think of different functions for the object and then respond according to how many were generated (1 to 5). For the *Read* task, subjects were to silently pronounce the word backward and rate the difficulty (1 = easy to 5 = hard). The study phase was subdivided into miniblocks of three consecutive trials, during which a single encoding task was performed. The miniblocks were employed to allow for efficient segregation of the hemodynamic responses according to task (also see McDuff et al., 2009), while not requiring long lags between individual study items (which would have significantly increased scanning time). Miniblocks began with a 4 s display of a task instruction (e.g., *Do ARTIST task*) and the response options, which remained on the screen throughout the miniblock. Each word appeared in the center of the screen for 2 s, and subjects were instructed to withhold their response until a response cue (*) appeared. Both the word and response cue remained on the screen for two seconds. Responses were made by pressing one of five keys mapped to the right hand. The second and third words of the miniblock were presented similarly and followed immediately by another miniblock. Each study block comprised 18 miniblocks (six per task) which were randomly ordered such that no task was completed twice consecutively. (An analysis of behavioral performance during the study phase is reported in the Supplemental Data.)

For the test phase, subjects were shown a series of intermixed old and new words and required to make one of five responses to each word (following Yonelinas et al., 2005). Subjects were to respond with their right thumb when they could remember specific details surrounding the word's presentation during the study phase (Remember). It was emphasized that subjects should give a Remember response if they remembered any details, regardless of whether the details were directly related to the study tasks or unrelated. The instructions also included a description of some examples of task-unrelated details, such as a personal thought elicited by a study item (e.g., something about your own dog in response to seeing the word *dog*) and an environmental stimulus co-occurring with an item (e.g., an unexpected background noise). If no study details were remembered, subjects used a four-point scale to rate their confidence that the word was either old or new. The right index through little fingers were mapped respectively to Sure Old, Unsure Old, Unsure New, and Sure New responses. Each test word was displayed centrally for three seconds, during which subjects were instructed to make their response. Responses outside the 3 s period were infrequent and not analyzed. Test words were followed by relatively long inter-item lags, during which a plus sign was centrally displayed, which helped to segregate the hemodynamic responses elicited by individual items. Each test block contained 48 trials with 5 s lags, 18 with 7 s lags, and 6 with 9 s lags, divided equally between old and new items.

Data Acquisition and Preparation

fMRI data were acquired with a Siemens Allegra 3T scanner at the Center for the Study of Brain, Mind, and Behavior at Princeton University. A T_1 -weighted anatomical volume (176 sagittal slices, 2 s TR, 4.38 ms TE, 1 mm³ voxels, 78° flip angle, and 256 mm² FOV) was acquired with an MP-RAGE sequence. Functional volumes consisted of T_2^* -weighted echoplanar images with blood oxygenation level dependent (BOLD) contrast and the following parameters: 2 s TR, 30 ms TE, 34 slices, 3.9 mm slice thickness, 3 mm² in-plane resolution, 75° flip angle, and 192 mm² FOV. The fMRI data were acquired in 7 separate blocks, with 152 volumes for each of 3 study blocks and 326 volumes for each of 4 test blocks. Five additional fMRI volumes collected at the beginning of each block permitted T_1 equilibration and were discarded before analysis. The onset of each study and test item coincided with the acquisition onset of an fMRI volume. The fMRI data were pre-processed using the AFNI software package (Cox, 1996; <http://afni.nimh.nih.gov/afni>). All volumes were spatially realigned to the first volume of the first study block, and the data in each volume were temporally shifted to the onset of the middle slice. Voxels exhibiting signal spikes were replaced via a temporal smoothing algorithm. Linear and quadratic trends were removed from each run to minimize the influence of scanner drift. The fMRI data were z scored separately for each voxel and

block. Notably, the fMRI data were neither spatially normalized nor smoothed prior to being used for the classification analyses.

For the classification analyses (see below), only a subset of the fMRI data were used—those volumes (TRs) determined as being associated with the study and test items, based on the lag in timing between stimulus events and the assumed resulting hemodynamic response. For the study phase, the hemodynamic lag was accounted for by convolving the onset of each study word with a synthetic hemodynamic response function (HRF; the gamma variant of AFNI's *waver*). Due to our use of miniblocks of study items assigned to a single task, this convolution produced a relatively dispersed (boxcar-like) HRF for each miniblock rather than three distinct item-specific HRFs. The convolved values at each time point (TR) were then normalized (from 0 to 1) across time. TRs with normalized values ≥ 0.5 were assigned to the corresponding (immediately presented) study task, whereas all other study phase TRs were excluded from the classification. With the first TR (hereafter TR 1) marking the onset of a miniblock's first word, the binarization procedure resulted in TRs 4 through 9 being assigned to the task completed during that miniblock. Given the relatively slow cycling through miniblocks, no study phase TR was assigned to more than one task. For the test phase, the fMRI volumes used in the classification corresponded to the TRs during which items onset (TR 1) followed by six subsequent TRs (2 through 7). Using a series of consecutive test phase TRs allowed us to assess classifier performance as a function of time.

For display purposes, each subject's anatomical data were normalized to a standard T_1 -weighted template (ICBM452; <http://www.loni.ucla.edu>) in Talairach space (Talairach and Tournoux, 1988). The resulting normalization parameters were also applied to the results of the individual-subject classification analyses (i.e., the importance and searchlight maps, as described below), which were resampled into 3 mm^3 voxels, in order to perform additional group-based analyses. The importance maps were smoothed with an 8 mm FWHM Gaussian kernel in order to create groupwise maps. The searchlight maps were left unsmoothed.

Pattern-Classification Analyses

Analyses of the fMRI data were performed with the Multi-Voxel Pattern Analysis toolbox (MVPA; Computational Memory Laboratory, Princeton, NJ; <http://www.csmbm.princeton.edu/mvpa>) and SPM5 (Wellcome Department of Imaging Neuroscience, London, UK; <http://www.fil.ion.ucl.ac.uk/spm>) in MATLAB (The MathWorks, Natick, MA). MVPA involves using neural network classifiers to determine how patterns expressed in multiple voxels of fMRI data relate to different experimental conditions. In the present implementation, classifiers were trained on study phase data and then validated on test phase data. The ability of a classifier to determine the prior encoding condition of a test item was used as the putative measure of reinstatement. The present study employed two types of classification (Whole-brain MVPA and Searchlight MVPA). Both types were conducted on an individual subject basis, with the reported results reflecting group-wise descriptors or further group analyses of the individual results.

Whole-Brain MVPA

The whole-brain MVPA procedure was similar to that used previously (McDuff et al., 2009; Polyn et al., 2005). The classifier consisted of a two layer (input and output; no hidden layer) feed-forward neural network, with full connections between input and output nodes. The input layer represented the fMRI data (one node per voxel) and the output layer corresponded to the task conditions (three nodes representing Artist, Function, and Read). A feature selection procedure (see below and Supplemental Data) was used to select the voxels to be included in the classification, regardless of how these voxels were distributed throughout the brain.

Training of the classifier began by initializing the input-output connection weights to random values between 0 and 1. Each training pattern of study phase fMRI data was then submitted to the classifier in random order. Classifier output for a given training pattern was determined by a sigmoid transfer function, producing values between 0 and 1. A cross-entropy function was used to calculate the classifier's prediction error following each training pattern, based on a comparison of the actual (true) and computed outputs. For example, the actual output values for a TR corresponding to the Artist task would be 1/0/0 for Artist/Function/Read. The classifier's weights were

updated with a conjugate gradient descent version of the backpropagation algorithm (for further discussion, see Bishop, 1995; Duda et al., 2001; Rumelhart et al., 1996). Training continued until either the mean prediction error across the three output nodes fell below .001 or there were 500 passes through all of the training TRs. Subsequently, classifier validation involved submitting each test phase TR of fMRI data to the trained network and noting the resulting (computed) values of the output nodes. To reduce the prediction error associated with randomly initializing the network weights, the classification was repeated 50 times for each subject, with a fresh randomization for each repetition. Results reflect the average across the 50 repetitions.

Classifier performance can be hindered by the inclusion of input data that exhibit excessive noise or are uninformative of the experimental conditions (for further discussion, see Mitchell et al., 2004, and Norman et al., 2006). To maximize performance we implemented a feature (voxel) selection procedure that restricted classifier input to only those voxels showing the largest differences among the three study tasks. Using an additional independent classifier based solely on the study phase data, the optimal number of voxels was determined to be 1000 (consistent with McDuff et al., 2009; see Supplemental Data). For each subject, voxel selection began by setting up a GLM (implemented in AFNI) that included a regressor for the convolved time course of each study task and nuisance regressors generated from spatial realignment. The F values from an ANOVA of the parameter estimates for the three tasks were then sorted, with the 1000 voxels exhibiting the largest values selected as input data.

To identify the voxels that were most influential in determining classifier output across subjects, we created importance maps for each subject by multiplying the average value of each input node by the three weights (post-training) connecting that node to the output layer. Voxels with positive values for both activity and weight were assigned positive importance values, voxels with negative activity and weight were assigned negative importance values, and voxels for which the activity and weight had opposite signs were assigned importance values of zero (McDuff et al., 2009; cf. Polyn et al., 2005). An across-subjects average map was created for each task, following spatial normalization and smoothing (see above).

Searchlight MVPA

The second type of classification followed an information-based searchlight approach (e.g., Kriegeskorte et al., 2006). For these analyses, the fMRI data were first divided into searchlights, consisting of all voxels falling within a sphere with a radius of 2 voxels. Each searchlight thus contained 33 voxels. A searchlight was centered on each voxel in a subject's brain, truncating those searchlights at the edge of the brain so as to exclude nonbrain voxels.

A separate classification was conducted for each searchlight. As in our previous classifications, the input layer of the classifier consisted of the fMRI data (one node for each of the 33 voxels) while the output layer corresponded to the encoding tasks. We found that the searchlight analysis ran too slowly when we used our standard backpropagation classification procedure, so we switched to using Gaussian Naive Bayes (GNB) classifiers for the searchlight analysis (Mitchell et al., 2004). GNB classification runs faster than backpropagation because it computes the input-output weights analytically (in contrast to backpropagation, which sets weights via an iterative error-correction procedure).

The task outputs generated for a given searchlight were assigned to its center voxel. Thus, the searchlight results constituted whole-brain maps of outputs for each of the three study tasks and for each TR during the test phase. To simplify the results, maps corresponding to TRs 3 through 5 for a given test item were averaged into a single map, and were then averaged according to experimental conditions. After spatiale normalization of the searchlight maps, they were imported into SPM5 for further groupwise analysis. All of the effects reported as significant survived a cluster-wise threshold of $p < 0.05$ (Worsley et al., 1996).

SUPPLEMENTAL DATA

Supplemental Data include six figures, one table, and supplemental text and can be found with this article online at [http://www.cell.com/neuron/supplemental/S0896-6273\(09\)00622-9](http://www.cell.com/neuron/supplemental/S0896-6273(09)00622-9).

ACKNOWLEDGMENTS

This work was supported by the National Institutes of Health (R01-MH072966 to M.D.R. and P50-MH062196 to K.A.N.).

Accepted: August 4, 2009

Published: September 9, 2009

REFERENCES

- Alvarez, P., and Squire, L.R. (1994). Memory consolidation and the medial temporal lobe: a simple network model. *Proc. Natl. Acad. Sci. USA* *91*, 7041–7045.
- Bishop, C. (1995). *Neural Networks for Pattern Recognition* (New York: Oxford University Press).
- Coltheart, M. (1981). The MRC psycholinguistic database. *Q. J. Exp. Psychol. A* *33*, 497–505.
- Cox, R.W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput. Biomed. Res.* *29*, 162–173.
- Damasio, A.R. (1989). Time-locked multiregional retroactivation: a systems-level proposal for the neural substrates of recall and recognition. *Cognition* *33*, 25–62.
- Davachi, L., Mitchell, J.P., and Wagner, A.D. (2003). Multiple routes to memory: distinct medial temporal lobe processes build item and source memories. *Proc. Natl. Acad. Sci. USA* *100*, 2157–2162.
- Donaldson, W. (1996). The role of decision processes in remembering and knowing. *Mem. Cognit.* *24*, 523–533.
- Duda, R., Hart, P., and Stork, D. (2001). *Pattern Classification, Second Edition* (New York: Wiley).
- Dunn, J.C. (2004). Remember-know: a matter of confidence. *Psychol. Rev.* *111*, 524–542.
- Düzel, E., Habib, R., Rotte, M., Guderian, S., Tulving, E., and Heinze, H.J. (2003). Human hippocampal and parahippocampal activity during visual associative recognition memory for spatial and nonspatial stimulus configurations. *J. Neurosci.* *23*, 9439–9444.
- Dzulkifli, M.A., and Wilding, E.L. (2005). Electrophysiological indices of strategic episodic retrieval processing. *Neuropsychologia* *43*, 1152–1162.
- Greenhouse, G.W., and Geisser, S. (1959). On methods in the analysis of repeated measures designs. *Psychometrika* *49*, 95–112.
- Haynes, J.D., and Rees, G. (2006). Decoding mental states from brain activity in humans. *Nat. Rev. Neurosci.* *7*, 523–534.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scand. J. Stat.* *6*, 65–70.
- Johnson, J.D., and Rugg, M.D. (2007). Recollection and the reinstatement of encoding-related cortical activity. *Cereb. Cortex* *17*, 2507–2515.
- Johnson, M.K., Hashtroudi, S., and Lindsay, D.S. (1993). Source monitoring. *Psychol. Bull.* *114*, 3–28.
- Johnson, M.K., Kounios, J., and Noldo, S.F. (1997). Electrophysiological brain activity and memory source monitoring. *Neuroreport* *8*, 1317–1320.
- Johnson, J.D., Minton, B.R., and Rugg, M.D. (2008). Content dependence of the electrophysiological correlates of recollection. *Neuroimage* *39*, 406–416.
- Kahn, I., Davachi, L., and Wagner, A.D. (2004). Functional-neuroanatomic correlates of recollection: implications for models of recognition memory. *J. Neurosci.* *24*, 4172–4180.
- Kriegeskorte, N., Goebel, R., and Bandettini, P. (2006). Information-based functional brain mapping. *Proc. Natl. Acad. Sci. USA* *103*, 3863–3868.
- Kucera, H., and Francis, W.N. (1967). *Computational Analysis of Present-Day American English* (Providence, RI: Brown University Press).
- Mandler, G. (1980). Recognizing: the judgment of previous occurrence. *Psychol. Rev.* *87*, 252–271.
- Maratos, E.J., Dolan, R.J., Morris, J.S., Henson, R.N.A., and Rugg, M.D. (2001). Neural activity associated with episodic memory for emotional context. *Neuropsychologia* *39*, 910–920.
- McClelland, J.L., McNaughton, B.L., and O'Reilly, R.C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* *102*, 419–457.
- McDuff, S.G., Frankel, H.C., and Norman, K.A. (2009). Multivoxel pattern analysis reveals increased memory targeting and reduced use of retrieved details during single-agenda source monitoring. *J. Neurosci.* *29*, 508–516.
- Mitchell, T., Hutchinson, R., Niculescu, S., Pereira, F., Wang, X., Just, M., and Newman, S. (2004). Learning to decode cognitive states from brain images. *Mach. Learn.* *57*, 145–175.
- Mur, M., Bandettini, P.A., and Kriegeskorte, N. (2009). Revealing representational content with pattern-information fMRI—an introductory guide. *Soc. Cogn. Affect. Neurosci.* *4*, 101–109.
- Norman, K.A., and O'Reilly, R.C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. *Psychol. Rev.* *110*, 611–646.
- Norman, K.A., Polyn, S.M., Detre, G.J., and Haxby, J.V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn. Sci.* *10*, 424–430.
- O'Reilly, R.C., and Rudy, J.W. (2001). Conjunctive representations in learning and memory: principles of cortical and hippocampal function. *Psychol. Rev.* *108*, 311–345.
- O'Reilly, R.C., and Norman, K.A. (2002). Hippocampal and neocortical contributions to memory: advances in the complementary learning systems framework. *Trends Cogn. Sci.* *6*, 505–510.
- Polyn, S.M., Natu, V.S., Cohen, J.D., and Norman, K.A. (2005). Category-specific cortical activity precedes retrieval during memory search. *Science* *310*, 1963–1966.
- Rolls, E.T. (2000). Hippocampo-cortical and cortico-cortical backprojections. *Hippocampus* *10*, 380–388.
- Rubin, D.C., and Greenberg, D.L. (1998). Visual memory-deficit amnesia: a distinct amnesic presentation and etiology. *Proc. Natl. Acad. Sci. USA* *95*, 5413–5416.
- Rugg, M.D., Johnson, J.D., Park, H., and Uncapher, M.R. (2008). Encoding-retrieval overlap in human episodic memory: a functional neuroimaging perspective. *Prog. Brain Res.* *169*, 339–352.
- Rumelhart, D., Durbin, R., Golden, R., and Chauvin, Y. (1996). Backpropagation: the basic theory. In *Backpropagation: Theory, Architectures, and Applications*, Y. Chauvin and D. Rumelhart, eds. (Mahwah, NJ: Erlbaum), pp. 1–34.
- Shastri, L. (2002). Episodic memory and cortico-hippocampal interactions. *Trends Cogn. Sci.* *6*, 162–168.
- Stark, C.E.L., and Okado, Y. (2003). Making memories without trying: Medial temporal lobe activity associated with incidental memory formation during recognition. *J. Neurosci.* *23*, 6748–6753.
- Stern, C.E., Corkin, S., Gonzalez, R.G., Guimaraes, A.R., Baker, J.R., Jennings, P.J., Carr, C.A., Sugiura, R.M., Vedantham, V., and Rosen, B.R. (1996). The hippocampal formation participates in novel picture encoding: Evidence from functional magnetic resonance imaging. *Proc. Natl. Acad. Sci. USA* *93*, 8660–8665.
- Talairach, J., and Tournoux, P. (1988). *Co-planar Stereotaxic Atlas of the Human Brain. 3-dimensional Proportional System: An Approach to Cerebral Imaging* (New York: Thieme Medical Publishers).
- Tulving, E. (1985). Memory and consciousness. *Can. Psychol.* *26*, 1–12.
- Wais, P.E., Mickes, L., and Wixted, J.T. (2008). Remember/know judgments probe degrees of recollection. *J. Cogn. Neurosci.* *20*, 400–405.
- Wheeler, M.E., and Buckner, R.L. (2004). Functional-anatomic correlates of remembering and knowing. *Neuroimage* *21*, 1337–1349.
- Wilson, M. (1988). Psycholinguistic database: machine-usable dictionary, version 2.00. *Behav. Res. Methods Instrum. Comput.* *20*, 6–10.

- Wixted, J.T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychol. Rev.* *114*, 152–176.
- Wixted, J.T., and Stretch, V. (2004). In defense of the signal detection interpretation of remember/know judgments. *Psychon. Bull. Rev.* *11*, 616–641.
- Woodruff, C.C., Johnson, J.D., Uncapher, M.R., and Rugg, M.D. (2005). Content-specificity of the neural correlates of recollection. *Neuropsychologia* *43*, 1022–1032.
- Worsley, K.J., Marrett, S., Neelin, P., Vandal, A.C., Friston, K.J., and Evans, A.C. (1996). A unified statistical approach for determining significant signals in images of cerebral activation. *Hum. Brain Mapp.* *4*, 58–73.
- Yonelinas, A.P. (2002). The nature of recollection and familiarity: a review of 30 years of research. *J. Mem. Lang.* *46*, 441–517.
- Yonelinas, A.P., Otten, L.J., Shaw, K.N., and Rugg, M.D. (2005). Separating the brain regions involved in recollection and familiarity in recognition memory. *J. Neurosci.* *25*, 3002–3008.