NeuroImage 127 (2016) 144–157

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

Decisions about the past are guided by reinstatement of specific memories in the hippocampus and perirhinal cortex

Michael L. Mack [a,b,*], Alison R. Preston [a,b,c,*]

[a] Department of Psychology, The University of Texas at Austin, Austin, TX 78712, United States
[b] Center for Learning and Memory, The University of Texas at Austin, Austin, TX 78712, United States
[c] Department of Neuroscience, The University of Texas at Austin, Austin, TX, 78712, United States

ARTICLE INFO

Article history:
Received 16 July 2015
Accepted 10 December 2015
Available online 15 December 2015

Keywords:
Decision making
Hippocampus
Medial temporal lobe
Memory
Perirhinal cortex
Computational modeling

ABSTRACT

When faced with a new challenge, we often reflect on related past experiences to guide our behavior. The ability to retrieve memories that overlap with current experience, a process known as pattern completion, is theorized as a critical function of the hippocampus. Although this view has influenced research for decades, there is little empirical support for hippocampal pattern completion to individual memory elements and its influence on behavior. We used pattern analysis of brain activity measured with functional magnetic resonance imaging to demonstrate that specific elements of past experiences are reinstated in the hippocampus, as well as perirhinal cortex (PRC), when making decisions about those experiences. Linking neural measures of specific memory reinstatement in the hippocampus and PRC to behavior with computational modeling revealed that reinstatement predicts the speed of memory-based decisions. Moreover, hippocampal activation during retrieval was selectively coupled to regions of occipito-temporal cortex that showed content-specific item reinstatement. These results provide evidence for hippocampal pattern completion and its role in the mechanisms of decision making.

Decisions often rely on information that is not immediately available. Whether it is ordering an entrée at a favorite restaurant or deciding among treatment options for a medical condition, many of our decisions are guided by what we recall from past experience. Most research on the neural systems of decision making have focused on perceptual decisions that depend on information directly or recently available to the sensory system (Heekeren et al., 2008). Although such studies have characterized how external sensory information is represented and interpreted by the brain to influence behavior (Bogacz et al., 2010; Bollimunta et al., 2012; Gold and Shadlen, 2007; Hanks et al., 2014; Heekeren et al., 2008; Nienborg and Cumming, 2009), the neural computations that support how we use internally-generated content from past experiences to guide decision making remain poorly understood.

It is well known that our ability to encode and remember past experiences depends critically on the medial temporal lobe (MTL) (Eichenbaum, 2004; Squire, 2004). The prominent theoretical view (Marr, 1971; McClelland et al., 1995; O'Reilly and Rudy, 2001) proposes that successful memory retrieval involves reinstatement of memory representations through a process known as pattern completion. Specifically, the hippocampus is thought to act as an auto-associative network that reinstates complete memory representations from partial or degraded input. It has been further proposed that the content of these reinstated memories serve as internally-generated evidence that guides subsequent behavior (Lisman and Grace, 2005; Norman and O'Reilly, 2003). In the present study, our goal was to test the theoretical prediction that during memory-based decision making, specific memory elements are reinstated in the human hippocampus and that this reinstatement predicts decisions.

Empirical evidence for hippocampal memory reinstatement and its influence on decision making has been established in human electrophysiological studies. This work has demonstrated that firing patterns of single neurons in medial temporal lobe selectively code for specific contents of memory (Heit et al., 1988) and that reinstatement of these firing patterns during retrieval is related to memory performance (Gelbard-Sagiv et al., 2008; Paz et al., 2010; Rutishauser et al., 2015). However, such empirical evidence of pattern completion in human functional magnetic resonance imaging (fMRI) research has been limited. Initial reports focused on differences in process rather than contents of memories, inferring pattern completion from a reduction in blood-oxygen-level dependent (BOLD) response to repeated presentations of objects and similar lures relative to initial encoding (Bakker et al., 2008; Lacy et al., 2011). Such an effect is consistent with pattern completion; however, it is unknown whether or not such reduced neural responding reflects reinstatement of specific memory elements.

* Corresponding authors at: Center for Learning and Memory, The University of Texas at Austin, 100 East 24th Street, Austin, TX 78712-0805, United States.
E-mail addresses: mack.michael@gmail.com (M.L. Mack), apreston@utexas.edu (A.R. Preston).

http://dx.doi.org/10.1016/j.neuroimage.2015.12.015
1053-8119/

fMRI studies have demonstrated that hippocampal activation patterns during spatial context encoding exhibit a bias towards existing spatial templates, a finding consistent with pattern completion processes (Stokes et al., 2015). Further, it has been shown that hippocampal activation patterns during retrieval differentiate between stored event representations (Chadwick et al., 2010, 2011) and are consistent with episode-specific templates of activation patterns recorded after encoding and retrieval (Wimber et al., 2015). However, these extant findings cannot speak to the central tenet of pattern completion: hippocampal activation patterns representing specific memory elements are themselves reinstated during decisions that rely on memory.

A second theoretical aspect of pattern completion with limited empirical support is the proposal that reinstated mnemonic information in the hippocampus serves as evidence for memory-based decision making (Lisman and Grace, 2005; Norman and O'Reilly, 2003). It has been demonstrated that successful memory decisions are associated with stronger reinstatement signatures in content-specific cortical regions (Bosch et al., 2014; Johnson et al., 2009; Kuhl and Chun, 2014; LaRocque et al., 2013; Polyn et al., 2005; Ritchey et al., 2012; Staresina et al., 2012; Wing et al., 2015). However, the behavioral link between memory decisions and hippocampal pattern completion has yet to be observed. Furthermore, the link between reinstatement and behavior observed in these studies is based on the indirect comparison of cortical reinstatement signatures during successful versus unsuccessful memory choices, thus limiting the sensitivity of evaluating the direct influence of reinstatement on decision making on a trial-by-trial basis. Recent neuroimaging evidence shows that the magnitude of hippocampal activation is linked to memory confidence judgments (Leiker and Johnson, 2015; Thakral et al., 2015) and the speed of memory decisions (Gordon et al., 2014), with greater retrieval-related activation for higher confidence and faster responses. Electrophysiological work in humans has further identified single cells within the hippocampus that exhibit firing patterns consistent with a graded representation of memory strength and confidence during recognition decisions (Rutishauser et al., 2015). However, these studies did not index hippocampal pattern completion per se. And, although the single cell finding is a compelling demonstration of hippocampal cells representing memory evidence, the kind of information carried by these cells was not item specific, and the memory-specific firing patterns were observed in response to perceptual information currently available to the participant. The theoretically important question remains as to whether or not internally generated hippocampal memory evidence is directly related to decisions about the specific mnemonic content.

Here, we sought direct evidence of pattern completion during memory-based decision making by indexing reinstatement of the specific components of individual memories from patterns of fMRI activation in the human hippocampus and surrounding cortices. Participants learned arbitrary associations between pictures of objects and famous faces and places before performing a delayed match-to-memory test (Chen et al., 2011), in which objects cued memory retrieval of the associated face or place (Fig. 1). We used neural pattern similarity (Kriegeskorte et al., 2008; Xue et al., 2010) to assess the reinstatement of specific people and places during memory retrieval in parahippocampal cortex (PHC) and perirhinal cortex (PRC), MTL cortical regions that have demonstrated content-specific coding (e.g., LaRocque et al., 2013; Liang et al., 2013; Staresina et al., 2012), and hippocampus (Fig. 1). Item-specific patterns during retrieval were compared to item-specific neural patterns recorded during a pre-exposure phase, in which participants viewed each item in isolation prior to associative encoding. This approach allowed us to go beyond category level decoding of memory reinstatement (Johnson et al., 2009; Polyn et al., 2005), a method that has failed to provide evidence for reinstatement of hippocampal representations (Diana et al., 2008; Gordon et al., 2014; cf. Kuhl and Chun, 2014), to assess reinstatement of specific memory elements.

Furthermore, we investigated the theoretical proposal that reinstated mnemonic information in hippocampus serves as evidence for memory-based decision making (Lisman and Grace, 2005; Norman and O'Reilly, 2003). To this end, we assessed the link between neural reinstatement of specific memory content in MTL subregions, including the hippocampus, and decisions about the memory probes with a mathematical model of decision making commonly used to assess how perceptual evidence impacts decisions (Ratcliff, 1978). Critically,

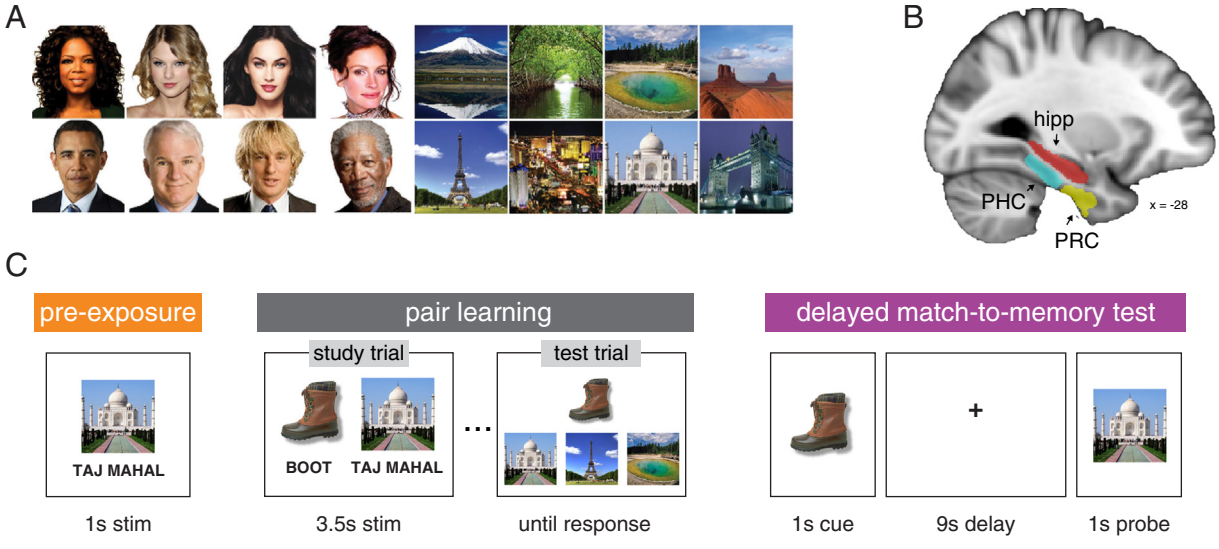

Fig. 1. Schematic of experimental task and depiction of brain regions of interest (ROIs). (A) Eight famous face and eight famous place images were used as stimuli in the experiment. (B) Analyses were conducted in MTL subregions: PHC (light blue), PRC (yellow), and the hippocampus (red). (C) The experiment consisted of three phases. In the pre-exposure phase, participants viewed images of faces, places, objects, and fixation points in isolation. During the pair learning phase, participants learned to criterion eighty paired associates consisting of a common real-world object and either a famous face or place through repetitions of first studying all pairs followed 3AFC tests on associate memory. Finally, participants were tested on each of the eighty pairs during a delayed match-to-memory (DMTM) test. Each test trial consisted of: an object cue (1 s), a delay period (9 s), a probe face or place (1 s) followed by an inter-trial interval (5-11 s, mean 7.5 s). The probe item was either the correct paired associate (a match) or a same-category foil drawn from another pair (a mismatch). Participants indicated match or mismatch status with a button response.

our model-based approach extends beyond correlation methods (e.g., Gordon et al., 2014) to test whether or not MTL neural reinstatement signatures are predictive of subsequent behavior. Typical approaches to relating brain measures to behavior are largely discriminative, focusing on differences in neural activation across known conditions. For example, studies that compare neural measures based on successful versus unsuccessful memory performance have revealed that cortical reinstatement is stronger for correct than incorrect memory decisions (e.g., Kuhl and Chun, 2014). The discriminative approach can reveal potentially important neural differences across conditions, but the resulting findings are ultimately correlational. In contrast, a generative approach attempts to describe the system of interest first and then evaluate the system under different conditions. Such an approach is considered generative because the modeled system can generate predictions for novel situations (Ng and Jordan, 2002). Here, we take a generative approach by using a modeling framework that makes explicit assumptions about how neural information available in the MTL influences memory decisions. By doing so, we provide a strong and direct test of the hypothesis that internally generated memory evidence in the hippocampus predicts mnemonic decision making on a trial-by-trial basis.

We also tested the key hypothesis that reinstatement in the hippocampus is linked to reinstatement in cortex (Ciocchi et al., 2015; Tanaka et al., 2014). It has been demonstrated that the magnitude of hippocampal BOLD activation during both encoding and retrieval scales with reinstatement signatures in neocortex (Gordon et al., 2014; Horner et al., 2015; Ritchey et al., 2012; Wing et al., 2015). Such findings suggest that the hippocampus plays an important role in the re-experiencing of mnemonic content by influencing cortical reinstatement. We extend this work by asking whether or not regions of the hippocampus and the surrounding cortex that demonstrate item-specific reinstatement are functionally coupled to regions of occipito-temporal cortex (OTC) that selectively reinstate specific mnemonic content. Such a functional link would suggest that not only is item-specific memory evidence reinstated in the hippocampus, but that the same underlying neural substrate is also communicating with regions of neocortex that demonstrate item-specific reinstatement.

Materials and methods

Participants

Twenty-five volunteers (15 females, mean age 20.2 years old, ranging from 18 to 29 years) participated in the experiment; one participant was excluded from analysis due to a failure in the button-response recording device and an additional participant was removed from further analysis due to substantial head motion. All subjects were right handed, had normal or corrected-to-normal vision, and were compensated $75 for participating.

Materials

Eight color images of famous faces (Oprah Winfrey, Taylor Swift, Meghan Fox, Julia Roberts, Barack Obama, Steve Martin, Owen Wilson, and Morgan Freeman) and famous places (Mount Fuji, Florida Everglades, Yosemite National Park, Monument Valley, Eiffel Tower, Las Vegas Strip, Taj Mahal, and London Bridge) as well as eighty-eight color images of everyday objects were used in the experiment (Fig. 1A). Face and object images consisted of a cropped face or object placed on a white background. All stimuli were sized to 300 × 300 pixels. Eight independent raters provided visual distinctiveness ratings for a set of 30 famous faces and places. The four most visually distinctive female and male faces and four most visually distinctive natural and man-made places were selected for the experiment.

Task procedure

After an initial screening and consent in accordance with the University of Texas Institutional Review Board, participants were instructed on the pre-exposure, associative learning, and delayed match to memory (DMTM) tasks. Participants then performed all three tasks in the MRI scanner by viewing visual stimuli back-projected onto a screen through a mirror attached onto the head coil. Foam pads were used to minimize head motion. Stimulus presentation and timing was performed using custom scripts written in Matlab (Mathworks) and Psychtoolbox (www.psychtoolbox.org) on an Apple Mac Pro computer running OS X 10.7.

Pre-exposure phase

During the pre-exposure, participants were instructed to respond as quickly as possible when a centrally-located fixation cross changed color to blue or yellow by pressing one of two buttons on a MRI-compatible button box. On each trial, a face, place, or object image was presented for 1 s along with a black fixation cross (Fig. 1C). After a random delay (250-750 ms) from image and fixation onset, the fixation cross changed color to blue or yellow. After 1 s had elapsed from the onset of the image and fixation cross, a blank screen was presented for 3-6 s. Participants were instructed to respond only to the fixation cross color change and that the image stimuli were irrelevant to the task. Trials were presented in a mixed design with blocks of eight trials consisting of stimuli from the same visual category (faces, places, and objects). Within a block, trials were randomly jittered with an average trial duration of 5.5 s and a range of 4-7 s. Each of the eight famous faces and places and eight objects were presented twice across a pre-exposure run. Five fixation blocks, each lasting 15 s, were presented interleaved with the stimulus blocks. Each run lasted 341 s, and participants completed four runs. The entire pre-exposure phase lasted approximately twenty-five minutes.

Associative learning phase

Each of the eight famous faces and places were paired with five objects for a total of eighty pairs. These pairings were randomly selected across participants. Participants performed repetitions of study and test blocks to learn the paired associates. Although participants performed the associative learning phase in the scanner, we did not collect fMRI data for this experimental phase. On each trial of a study block, a face-object or place-object stimulus pair was presented for 3.5 s followed by 0.5 s of fixation (Fig. 1C). The stimulus images were presented at the same time as the object image. The object image appeared 8.1° to the left of the screen center, and the face/place image 8.1° to the right of the screen center. Labels for the faces/places and objects were presented in black, 32-point, Helvetica text below the images. Each image subtended 7.3° × 7.3° of visual space. Participants were instructed to view and attempt to remember the pairs. All 80 pairs were shown in random order during a study block.

After study, participants were tested on the eighty paired associates. On each test trial, an object stimulus from one of the pairs was presented above three faces or places, with one face/place being the correct paired associate and two faces/places from different pairs serving as lures. Lures were always from the same category as the correct face or place. Participants responded by pressing one of three buttons on a MRI button box to indicate which of the three face/place stimuli was associated with the object. Test trials were self-paced with a 1 s feedback screen that presented the object along with the correct face/place. Test trials were separated by a 500 ms fixation. All 80 pairs were tested in random order. After each test block, the average accuracy was calculated across the test trials for that block. If the test accuracy was greater than or equal to 85%, participants moved onto the delayed match to memory task phase. Otherwise, participants repeated the study and test blocks with different randomized trial orders.

Delayed match to memory (DMTM) task

The test phase consisted of trials of a delayed match to memory task. In contrast to the widely-used delayed match to sample task, in which a stimulus is presented and held in working memory through a delay period before a match decision is made to a probe stimulus, the delayed match to memory task requires cued retrieval from long-term memory during the delay period. On each trial of the DMTM phase, an object from one of the eighty pairs encoded during the associative learning phase was presented as a memory cue for 1 s followed by a 9 s blank screen. After the 9 s delay, a memory probe was presented that was either the correct face/place paired associate (a match trial) or a different face/place associated with a different object (a mismatch trial). The probe stimulus was presented for 1 s before presentation of a 3 s fixation (Fig. 1C). Participants were instructed to respond whether or not the probe stimulus was an associative match or mismatch to the object cue by pressing one of two buttons on a MRI compatible button box. Participants were also instructed to respond within 1 s of the probe onset. A strict response deadline was used to encourage memory retrieval of the associated face/place during the delay period. Each run of the DMTM phase consisted of sixteen trials. Across these sixteen trials, each of the eight famous faces and places was a retrieval target, half of which were match trials and the other half mismatch trials. Across five runs all objects from the eighty learned pairs were presented as a cue. Null fixation events with durations of 2-8 s sampled from a truncated exponential distribution for a total of 72 s were randomly intermixed between trials throughout each run to maximize the efficiency of the trial order. Each run lasted 294 s with the entire memory test phase lasting approximately thirty minutes.

MRI data acquisition

Whole-brain imaging data were acquired on a 3.0 T Siemens Skyra system at the University of Texas at Austin Imaging Research Center. A high-resolution T1-weighted MPRAGE structural volume (TR = 1.9 s, TE = 2.43 ms, flip angle = 9°, FOV = 256 mm, matrix = 256 × 256, voxel dimensions = 1 mm isotropic) was acquired for coregistration and parcellation. Functional images were acquired using a T2*-weighted multiband accelerated EPI pulse sequence (TR = 1 s, TE = 30 ms, flip angle = 62°, FOV = 220 mm, matrix = 110 × 110, slice thickness = 2 mm, multiband factor = 4) allowing for whole brain coverage with 2 mm isotropic voxels.

MRI data preprocessing and statistical analysis

MRI data were preprocessed and analyzed using FSL 5.0 (Smith et al., 2004) and custom Python routines. Functional images were realigned to the first volume of the fourth run from the preprocessing phase to correct for motion, spatially smoothed using a 4 mm full-width-half-maximum Gaussian kernel, high-pass filtered (128 s), and detrended to remove linear trends within each run. Functional images were registered to the MPRAGE structural volume using Advanced Normalization Tools (ANTS, version 1.9; Avants et al., 2011). All analyses were performed in the native space of each participant.

Regions of interest

The hippocampus and subregions of MTL cortex (PRC and PHC) were delineated by hand on the 1 mm Montreal Neurological Institute (MNI) template brain (Fig. 1B) by an expert trained to segment MTL and hippocampus. Given the established functional differences in lateral versus medial subregions of entorhinal cortex (ERC; e.g., Knierim et al., 2014) and the limited resolution of fMRI, we excluded ERC from our analyses. The segmentations were validated by two other expert raters and have been used in previous published studies (Schlichting and Preston, 2014; Schlichting et al., 2015). PRC was defined as the lateral and medial wall of the collateral sulcus and a 3 mm wide transition zone between PHC and PRC was excluded from both regions. Hippocampal, PHC, and PRC ROIs were reverse-normalized to each individual's functional space using ANTS. Specifically, a nonlinear transformation was calculated from the MNI template brain to each participant's T1-weighted MPRAGE volume. This warp was then concatenated with the MPRAGE to functional space transformation calculated using ANTS. Finally, the concatenated transformation was applied to the anatomical MTL ROIs to move each ROI into each participant's functional space. See the Supplemental Materials for methods and findings from an exploratory analyses of hippocampal subfields $DG/CA_{2,3}$ and CA_1.

A region of interest including occipito-temporal cortex (OTC) was constructed from an automated cortical segmentation performed with Freesurfer, version 5.3 (Fischl et al., 2002). Specifically, the Freesurfer cortical segmentation was performed on each participant's T1-weighted MPRAGE volume and all cortical regions contained within occipito-temporal cortex, but excluding medial temporal lobe cortex, were added to an OTC ROI that was specific to each participant. The OTC ROIs was transformed to each participant's functional space using ANTS.

Pattern classification analysis

The goal of the classification analysis was to assess the extent to which category-level information (face vs. place) was reinstated in different MTL subregions during memory retrieval (Fig. 3). Pattern classification analyses were implemented using PyMVPA (Hanke et al., 2009) and custom Python routines. All classification analyses were conducted on preprocessed and spatially smoothed functional data. The decision to perform multivariate analyses on spatially smoothed data is consistent with recent studies employing MVPA in the MTL (Kuhl and Chun, 2014; e.g., Schapiro et al., 2012) and demonstrations that smoothing does not result in information loss (Kamitani and Sawahata, 2010; Op de Beeck, 2010). A linear support vector machine (SVM) classifier was trained to discriminate neural activation patterns of faces, places, objects, and fixation acquired during the pre-exposure phase. Each trial in the pre-exposure phase was labeled according to the stimulus category that was presented during the trial. To account for the hemodynamic lag, the fMRI data was time shifted by 4 s.

First, the classifier was evaluated for each of the MTL ROIs (PRC, PHC, and hippocampus) with a leave-one-out 4-fold cross-validation using the pre-exposure phase's four functional runs. Classifiers were trained to predict face, scene, object, and fixation blocks during the pre-exposure phase. All four categories were included in the classifier to isolate the four sources of potential information in neural patterns during the DMTM phase. By including both objects and fixation in the classifier, the delay period information associated with perceptual processing of the cue object and the fixation point was accounted for in the classifier evidences for those two categories. With this experimentally irrelevant variance accounted for, the classifier evidences for face and scene categories during the DMTM delay period was less biased and more accurately reflected category information reinstated through memory retrieval. Classifier performance during the pre-exposure phase was assessed by taking the average accuracy across the cross-validation folds within each participant, then evaluating whether or not across the participant's classification accuracy was significantly greater than chance levels (25% chance accuracy with four classes) according to a Bayesian Estimation Supersedes the *t*-Test procedure (BEST; Kruschke, 2012).

BEST is a hierarchical Bayesian version of the traditional *t*-test that estimates the *t*-test statistic through a Markov-Chain Monte Carlo (MCMC) sampling method. The BEST procedure is a generalization of the Student t-test such that the specific form of the *t* distribution (defined by the mean, standard deviation, and shape parameter) used by the test is estimated from the data through MCMC methods. Instead of assuming normally-distributed data (i.e., setting the shape parameter

equal to ∞) as is the case with traditional *t*-tests, the MCMC method estimates all three parameters of the *t* distribution defined by the data and provides a valid test between groups of data even when the data is not normally distributed. As the name of the test implies, the BEST procedure is a generalization of the traditional *t*-test and is a statistically superior method for evaluating means from group data (Kruschke, 2012). All BEST tests conducted in the current work consisted of 10,000 MCMC samples with the first 1000 samples discarded for burn-in. A BEST analysis provides a distribution of credible values of the test statistic summarized in the form of 95% highest density interval (HDI) that is akin to a 95% confidence interval from traditional t-test methods. Although the results of a BEST analysis are best summarized by the 95% HDI, we also calculated a *P* statistic that corresponds to the proportion of the sampling posterior probability distribution that falls below each test's chance level (i.e., *P* is an analog to a *p* value from traditional null hypothesis testing methods). Additionally, all reported *P* values are corrected to control for false discovery rate with $\alpha = 0.05$. All three ROIs showed greater than chance accuracy (all $P < 0.0001$) at decoding the four stimulus classes during the pre-exposure phase (PHC: mean = 32.9%, HDI = [31.5%, 34.3%]; PRC: mean = 28.7%, HDI = [27.2%, 29.7%]; hippocampus: mean = 27.3%, HDI = [26.2%, 28.3%]).

To further test the reliability of the pre-exposure MVPA classifiers, we ran randomization tests on all of the classifiers to empirically define chance levels. To perform these tests, we reran the classification analysis, but shuffled labels on the test data in each of the cross validation folds. This analysis was repeated 1000 times for each participant to create a null distribution of classifier accuracies. For each participant and each ROI, a *p* value was calculated that corresponded to the proportion of the null distribution accuracies that were greater than the observed accuracy. These *p* values were converted to *z* values with the inverse cumulative normal distribution and evaluated across participants with a one-sample BEST test. Both PHC and PRC classifiers showed significantly greater accuracies than a baseline z value of 1.65 (corresponding to a *p* value of 0.05; PHC: mean = 3.475, HDI = [2.562, 4.498], $P < 0.0001$; PRC: mean = 2.945, HDI = [1.954, 3.944], $P = 0.0074$). In contrast, accuracy of the hippocampus classifier across participants was not significantly greater than the randomization test null distributions (mean = 2.166, HDI = [1.188, 2.962], $P = 0.159$). Thus, whereas classifiers trained on pre-exposure data in PHC and PRC accurately classified above empirically defined chance levels, hippocampal classifiers did not. As such, subsequent classification analyses were restricted to PHC and PRC.

After assessing classifier performance during the pre-exposure phase, a SVM classifier was trained on all data from the pre-exposure phase for each ROI and used to predict evidence for face and place information during the delay period of the DMTM trials. This analysis was limited to correct trials. To account for the hemodynamic lag, the trained classifier was applied to data acquired at time points 4-7 s during the delay period of the DMTM trials (i.e., the dataset was time shifted by 4 s). This range was selected a priori to capture the peak of the hemodynamic response to the delay period. Probability estimates for the face and place labels were extracted from the SVM classifier and used to index the amount of classifier evidence for the two categories at each time point during the tested range. Each evidence estimate was averaged across the time points within a trial, relabeled as correct versus incorrect category according to the retrieval category of the trial, and averaged across memory test trials for each participant (Fig. 3). Finally, the difference in classifier evidence for correct versus incorrect category was evaluated across participants with a paired *t*-test based on the BEST procedure (Kruschke, 2012), with separate tests for memory test trials with face and place retrieval targets. For these particular tests, a mean paired difference was considered significant if the 95% HDI did not include zero. We have included a *P* statistic in the Results section that corresponds to the proportion of the sampling posterior probability distribution of mean differences that fell below zero. Finally, regional interactions between regions demonstrating significant category reinstatement were assessed with an analysis of variance that compared factors of region, retrieval category, and classifier category evidence.

Multivariate pattern similarity

The goal of the similarity analysis was to assess the extent that item-level information was reinstated in different MTL subregions during memory retrieval (Fig. 4). In contrast to classification techniques that are used to decode activation patterns associated with relatively small number of stimulus classes or conditions, pattern similarity methods allow one to evaluate activation patterns at the level of single events or stimuli. For example, instead of training a classifier to dissociate activation patterns for places and faces, as is the case with classification methods, pattern similarity methods can measure the degree that an activation pattern associated with a single trial of retrieving the Taj Mahal is similar to activation patterns associated with viewing the Taj Mahal versus other places. Pattern similarity methods offer a greater sensitivity in detecting between individual experimental events than is afforded by multivariate classification.

Pattern similarity analyses were implemented using PyMVPA (Hanke et al., 2009) and custom Python routines. First, activation patterns for visual and memory-related responses in each of the ROIs for each of the eight famous faces and places were estimated using an event-specific univariate general linear model (GLM) approach employing the least squares single (LSS) method (Mumford et al., 2012; Xue et al., 2010). In contrast to the classification approach that leverages the variance in neural patterns to learn voxel weights that best discriminate conditions, pattern similarity analyses require stable estimates of neural representations for the conditions of interest. In the current study, the condition of interest was at the level of specific items. Thus, we took a GLM approach to model stable estimates of visual and memory-related neural patterns for each of the eight famous faces and places. For each pre-exposure run, separate GLMs were conducted for each of the eight faces and places, such that each GLM included one regressor for the face or place modeled as 1 s boxcar convolved with a canonical hemodynamic response function (HRF), separate regressors for the remaining events collapsed according to the event categories (e.g., face, place, object, and fixation) that were also convolved with a HRF, and nuisance regressors that included six head motion parameters. These individual GLMs resulted in voxelwise parameter estimates to each of the stimuli. The resulting beta values consisted of four visually-evoked activation patterns across the runs for each of the eight faces and places. A similar GLM approach was conducted on the DMTM functional data to extract memory retrieval-based activation patterns for the faces and places. Specifically, for each DMTM functional run, separate GLMs were conducted for the individual trials such that each GLM included a single regressor for the current trial modeled as a 9 s boxcar convolved with a canonical HRF, separate regressors that combined the remaining events according to category and match condition (face match, face mismatch, place match, and place mismatch), and nuisance regressors that included six head motion parameters. These GLMs resulted in voxelwise parameter estimates of activation related to retrieving each of the faces and places. The resulting beta values consisted of five memory-retrieval activation patterns across the runs for each of the eight faces and places. In total, we estimated four visually-evoked and five retrieval-based activation activation patterns for each of the eight faces and places.

The pre-exposure and DMTM test activation patterns were used to characterize the degree of item-specific information reinstated within the ROIs during memory retrieval (Fig. 4). Pearson correlations were computed to assess the similarity between DMTM test patterns of retrieving an item and pre-exposure patterns of viewing that same item (e.g., retrieving Taj Mahal and viewing Taj Mahal, *self similarity*) or viewing other items within the same category (retrieving Taj Mahal and viewing Eiffel Tower, *category similarity*). Self similarity and

category similarity values were calculated for every DMTM test trial for each participant in the hippocampus, PRC, and PHC. For the hippocampal regions demonstrating an item reinstatement effect, we performed follow up exploratory analyses in CA subfields (see Supplemental Materials). Correlation coefficients were Fisher transformed to more closely conform to the assumptions of normality underlying standard statistical tests.

To assess the degree of item-level reinstatement within each ROI, we compared average self versus category similarity across participants for face and scene retrieval trials separately. Only correct trials with response times faster than or equal to 1000 ms were included in the participant averages. Given the visual nature of the materials and the implementation of the analysis, which emphasizes visual similarity, we expected that right hemisphere might be more likely to demonstrate item reinstatement effects (Glosser et al., 1995; Golby et al., 2001). We therefore first assessed differences in the similarity measures between left and right hemispheres for each ROI by conducting an analysis of variance that included factors of hemisphere (left vs. right), retrieval target (face vs. place), and similarity (self vs. category). If an ROI showed a significant interaction that included hemisphere, subsequent analyses analyzed the ROI separated by hemisphere; otherwise, similarity measures were collapsed across hemisphere. Pairwise comparisons between self and category similarity were conducted with BEST *t*-tests. Reported BEST *P* values are corrected to control for false discovery rate with an alpha = 0.05. Regional interactions between MTL regions showing item reinstatement were evaluated with an analysis of variance that compared factors of region, retrieval category, and similarity level. Any item reinstatement effects found in the hippocampus were followed up with exploratory pattern similarity analyses in subfield CA_1 and the combined subfields $DG/CA_{2,3}$ (see Supplemental Materials).

Additionally, we identified regions in occipito-temporal cortices (OTC) that showed item-level reinstatement with a searchlight analysis extension of the pattern similarity approach. This analysis was performed as part of the connectivity analysis described in the next section. Specifically, we targeted OTC regions that selectively coded for items from one category (face or place) more than the other to assess the functional connectivity between stimulus representation regions and MTL regions associated with item reinstatement. We used a searchlight approach due to the relatively large volume of OTC. The spatial sensitivity of a searchlight approach allowed us to identify only those regions across all of OTC that reinstated item information during memory retrieval. The same pattern similarity analysis performed within the anatomically-defined MTL ROIs described previously was employed for the searchlight analysis with the following exceptions. For each searchlight sphere (radius = 3 voxels), an item reinstatement index was calculated by subtracting the category similarity from the item similarity. This item reinstatement index was compared to a null distribution of reinstatement indices calculated from randomly shuffled trial orderings. This permutation test was conducted for each searchlight sphere resulting in a probability map.

The searchlight analysis was conducted separately for face and place item reinstatement measures and was restricted to occipito-temporal cortex as defined by Freesurfer anatomical parcellations. The resulting probability maps were converted to *z*-scores. For group analysis, the *z*-scored maps were normalized to MNI space and voxelwise nonparametric permutation testing (5000 permutations) was performed using FSL Randomise (Winkler et al., 2014) to compare face and place item reinstatement. The resulting statistical maps were voxelwise thresholded at $p = 0.01$ and cluster corrected at $p = 0.05$ which corresponded to a cluster extent threshold of greater than 29 voxels as determined by AFNI 3dClustSim (estimated smoothness of FWHM = 6.27). Significant face and place item reinstatement clusters are reported in Table 1 and were used to create ROI masks for the connectivity analysis described next (Fig. 5A).

Functional connectivity analysis

We assessed the functional coupling (Rissman et al., 2004) between MTL regions showing item-level reinstatement and item-reinstatement regions in OTC identified in the pattern similarity searchlight analysis described in the previous section (Fig. 5). For each trial, average delay-period activation within these face and place reinstatement regions, as well as the MTL ROIs exhibiting item reinstatement (right hippocampus and left PRC, see Supplemental Materials for right CA_1), were extracted from the event-specific univariate GLM described in the pattern similarity analysis section (Fig. 5A). For each participant, we estimated the degree of connectivity between the MTL ROIs and OTC face and place reinstatement regions separately for face and place retrieval trials using robust regression (Fig. 5B). Connectivity differences across different retrieval conditions (see Fig. 5C) were assessed with a repeated measure analysis of variance of retrieval type (face vs. place) and OTC reinstatement region (face vs. place). In particular, we looked for a significant interaction as evidence of connectivity modulated by retrieval type. We also performed paired comparisons within each MTL ROI for connectivity with OTC face and place reinstatement regions for each retrieval condition separately using the BEST methods described previously.

Diffusion modeling analysis

One key theoretical claim of pattern completion is that the mnemonic content reinstated in the hippocampus serves as evidence for memory based decisions. We sought to evaluate this prediction by linking our pattern similarity measure of item reinstatement to participants' match/mismatch decisions with a computational model of decision making. Specifically, we fit participants' choice behavior in the DMTM task with variants of the drift diffusion model (Ratcliff, 1978) that were differently informed by the neural item reinstatement measure. By doing so, we evaluated whether or not the fidelity of internally generated item reinstatement predicts how quickly responses to the retrieved memory are made on a trial-by-trial basis. This analysis extends beyond correlation methods (e.g., Gordon et al., 2014) by assessing the role of item reinstatement within a computational framework of memory-based decision making.

The DDM posits that decisions are made by sequentially sampling noisy evidence. Over time, this evidence accumulates towards a threshold. Once the threshold is reached, the response associated with the threshold is made. With this conceptualization, the DDM is able to account for both response probabilities and response time distributions. The DDM is based on three core parameters: drift rate v – the rate of evidence accumulation, threshold a – the distance between the starting point of the evidence and the decision thresholds, and non-decision time t – the time required for perceptual encoding and motor response execution. These parameters uniquely affect predictions for response probabilities and response time distributions. Typical applications of the DDM seek to decompose effects of experimental conditions on choice behavior onto one or more of the DDM parameters to gain insight into the latent mechanisms of decision making (Ratcliff and Rouder, 1998). Here, we used the DDM to investigate if and how the neural measure of item reinstatement predicts, for every trial, response choice and speed.

The DDM analysis was implemented with the Hierarchical Drift Diffusion Model toolbox (HDDM version 0.5.3; Wiecki et al., 2013). HDDM performs hierarchical Bayesian parameter estimation for the DDM using Markov-Chain Monte Carlo (MCMC) sampling methods. By doing so, the HDDM allows for simultaneous estimation of participant and group parameters such that participant parameters are sampled from group distributions, multiple parameters to be estimated for different conditions, and the estimation of trial-by-trial effects from external variables (e.g., neural measures) on different DDM parameters.

We fit a family of DDM variants to participant behavior, with different model variants representing different hypotheses about how neural item reinstatement influences decision making (Fig. 6). First, we calculated a trial-by-trial item reinstatement score by subtracting the category from

the item similarity values as computed in the neural pattern similarity analyses described previously. To remove the potential confound of trial-by-trial BOLD activation on choice behavior within each ROI, overall BOLD response within each ROI was computed for each trial and regressed out of the trial-by-trial item reinstatement score. We then used this modified item reinstatement score in the DDM analyses. Specifically, in separate models, we estimated the relationship between trial-by-trial item reinstatement in MTL regions showing evidence for item reinstatement, as well as OTC regions identified in the searchlight analysis, and the three core DDM parameters. This relationship was modeled as a linear effect between item reinstatement and a DDM parameter. Thus, the result of one model fit was two regression coefficients that describe the linear relationship between the neural measure of item reinstatement and the DDM parameter for match and mismatch trials.

For example, fitting a model with item reinstatement linked to drift rate (Fig. 6A) results in linear regression coefficients that characterize the effect of item reinstatement on drift rate separately for match and mismatch trials. Positive valued coefficients would indicate that a higher item reinstatement is associated with a higher drift rate. Negative valued coefficients would indicate that higher item reinstatement is associated with a lower drift rate. Because parameter estimation in the HDDM is implemented in a hierarchical Bayesian framework with MCMC sampling, a distribution of regression coefficients is estimated for any one model fit (Fig. 6B). To assess the linear effect of item reinstatement on the model parameters, we calculated a *P* statistic that quantified the proportion of the regression coefficient distribution that was less than zero. *P* values less than 0.05 were considered statistically significant. Separate models were fit to behavior associated with face and place retrieval trials, and all three DDM parameters (*v*, *a*, *t*) were free to vary across match and mismatch conditions.

All models were estimated with 30,000 MCMC samples with the first 10,000 samples discarded for burn-in. We assessed model fits with two analyses. First, model convergence was evaluated with the Gelman-Rubin *R-hat* statistic computed over three independently estimated chains of each model. All parameters in all models had a *R-hat* less than 1.1, suggesting the chains of MCMC sampling reached convergence. Second, we performed model comparison by calculating a deviance information criterion (DIC) measure for each model. Critically, we compared the item reinstatement models to baseline DDM models that did not include effects of item reinstatement. An item reinstatement model was considered as representing a meaningful relationship between behavior and item reinstatement if the model provided a better fit (smaller DIC value) than the corresponding baseline model (Table 2). If more than one item reinstatement model provided a better fit than the baseline model, the item reinstatement model with the smaller DIC value was selected as the winning model.

Results

Behavioral performance

For the pair encoding phase, learning performance was defined as the average number of correct responses during each test block. A criterion of 85% accuracy was used to determine when to move on to the DMTM phase. Of the 23 participants, 10 reached the learning criterion after one study-test repetition, 11 after two study-test repetitions, and 2 after three study-test repetitions.

For the DMTM phase, memory performance in each of the four conditions (face match, face mismatch, place match, and place mismatch) was measured with the mean probability of a correct response and the median response times for correct responses (Fig. 2). Trials were excluded from all analyses if no response was provided (1.1% of all trials) or if the trial response time was greater than 1000 ms (5.1% of all trials). An analysis of variance of probability of correct responses that included the factors of probe match (match vs. mismatch) and retrieval category (face vs. place) revealed a significant effect of retrieval category ($F_{1,22} = 4.34$, $p = 0.049$, $\eta_p^2 = 0.165$), but no effect of probe match ($F_{1,22} = 0.576$, $p = 0.456$, $\eta_p^2 = 0.026$) nor an interaction ($F_{1,22} = 0.167$, $p = 0.687$, $\eta_p^2 = 0.008$). A similar analysis of variance conducted on response times revealed significant main effects of probe match ($F_{1,22} = 15.793$, $p = 0.0006$, $\eta_p^2 = 0.418$) and retrieval category ($F_{1,22} = 78.661$, $p = 1 \times 10^{-8}$, $\eta_p^2 = 0.781$), but no interaction ($F_{1,22} = 0.197$, $p = 0.661$, $\eta_p^2 = 0.009$). Thus, participants were overall more accurate and faster in responding to associative pairs that included faces relative to places, as well as faster to respond when the probe item matched relative to mismatched the face or place paired with the cue object.

Category reinstatement in the MTL

We first targeted evidence of category reinstatement with a commonly-used decoding approach (e.g., Kuhl and Chun, 2014; Zeithamova et al., 2012). Separate classifiers were trained on fMRI data from the pre-exposure phase for each bilateral MTL ROI (PHC, PRC, and hippocampus) to dissociate neural patterns for the four stimulus categories: objects, faces, places, and fixation. Only classifiers trained on PHC and PRC pre-exposure data showed above chance classification. Hippocampus did not differentiate category information during pre-exposure, thus this region was excluded from subsequent classification analyses for the DMTM test phase. The trained PHC and PRC classifiers were then applied to data recorded during the DMTM test phase to predict evidence for face and place information during the delay period of the DMTM trials (Fig. 3A). This analysis was restricted to correct DMTM trials. There was a difference in category reinstatement across MTL cortex as revealed by an analysis of variance that showed a significant interaction of region × stimulus category × retrieval category ($F_{1,22} = 29.18$, $p = 0.00002$, $\eta_p^2 = 0.57$). Within PRC (Fig. 3B), we found category-specific evidence of face reinstatement ($\mu_{diff} = 0.0162$, 95% HDI = [0.0051, 0.0281], $P = 0.002$), but not place reinstatement ($\mu_{diff} = 0.00158$, 95% HDI = [−0.012, 0.015], $P = 0.41$). Within PHC, we found evidence of place reinstatement ($\mu_{diff} = 0.0414$, 95% HDI = [0.0197, 0.0653], $P < 0.0001$), and a trending effect for face reinstatement ($\mu_{diff} = 0.011$, 95% HDI = [−0.0014, 0.024], $P = 0.076$).

Item reinstatement in the MTL

We next targeted the specific item-level contents of memory by comparing memory-based activation patterns within the hippocampus and surrounding cortex to visually-evoked activation patterns with neural pattern similarity analyses. Specifically, the degree of item-specific reinstatement was measured by correlating activation patterns present during retrieval of a specific face or place in DMTM trials with the corresponding activation patterns present when viewing that face or place prior to associative learning (Fig. 4A). This *self similarity* measure was compared to a *category similarity* measure defined as the

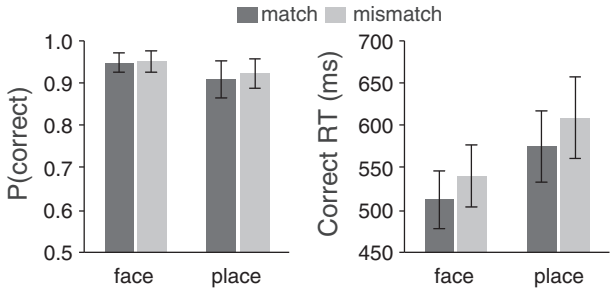

Fig. 2. Behavioral performance in the delayed match-to-memory (DMTM) test. Average probability correct (left plot) and median correct response times (right plot) are depicted for face and place retrieval trials where the probe was either a match (dark) or mismatch (light). Error bars represent 95% confidence intervals of the mean.

correlation between memory retrieval patterns for the faces and places and pre-exposure patterns for other items within the same stimulus class.

First, we assessed differences in the similarity measures between left and right hemispheres for each ROI by conducting an analysis of variance that included factors of hemisphere (left vs. right), retrieval target (face vs. place), and similarity (self vs. category). Given the visual nature of the experimental materials and the established laterality effects for stimulus modality in MTL (Glosser et al., 1995; Golby et al., 2001), we expected differences across hemispheres. If an ROI showed a significant interaction that included hemisphere, subsequent analyses analyzed the ROI separated by hemisphere; otherwise, similarity measures were collapsed across hemisphere. Both PRC and hippocampus showed significant interactions with hemisphere (PRC: hemisphere × target × similarity - $F_{1,22} = 9.87$, $p = 0.0044$, $\eta_p^2 = 0.291$; hippocampus: hemisphere × target - $F_{1,22} = 7.42$, $p = 0.0012$, $\eta_p^2 = 0.25$), thus corresponding ROIs in left and right hemispheres for these regions were analyzed separately.

Within right hippocampus (Fig. 4C), we found evidence of item-specific reinstatement as indicated by a main effect of item versus category similarity ($F_{1,22} = 6.99$, $p = 0.015$, $\eta_p^2 = 0.241$). Paired BEST comparisons revealed this main effect was driven by item reinstatement for places ($\mu_{diff} = 0.0085$, 95% HDI = [0.0002, 0.017], $P = 0.024$), but not faces ($\mu_{diff} = 0.0044$, 95% HDI = [−0.006, 0.0143], $P = 0.184$). Left hippocampus showed no significant item reinstatement effects ($p > 0.16$). In MTL cortex, only left PRC showed evidence of item-specific reinstatement as revealed by an interaction of similarity and retrieval target ($F_{1,22} = 5.41$, $p = 0.029$, $\eta_p^2 = 0.184$). Paired BEST comparisons indicated significant reinstatement for faces ($\mu_{diff} = 0.0099$, 95% HDI = [0.0002, 0.02], $P = 0.022$), but not places ($\mu_{diff} = -0.001$, 95% HDI = [−0.0112, 0.0092], $P = 0.86$). An analysis of variance assessing regional differences in reinstatement between right hippocampus and left PRC showed no significant interactions with region (all Fs < 1). Right PRC and bilateral PHC did not show evidence of item reinstatement (all Fs < 1). See the Supplemental Materials for an exploratory item reinstatement analysis of CA subfields in right hippocampus.

MTL connectivity with item reinstatement regions in occipito-temporal cortex

We next examined the key theoretical component of pattern completion that the hippocampus guides reinstatement in behaviorally-relevant cortical regions (Marr, 1971). We used a functional connectivity approach (Rissman et al., 2004) to examine the coupling between hippocampal processes and task-related cortical regions during retrieval. We identified regions throughout occipito-temporal cortex (OTC)

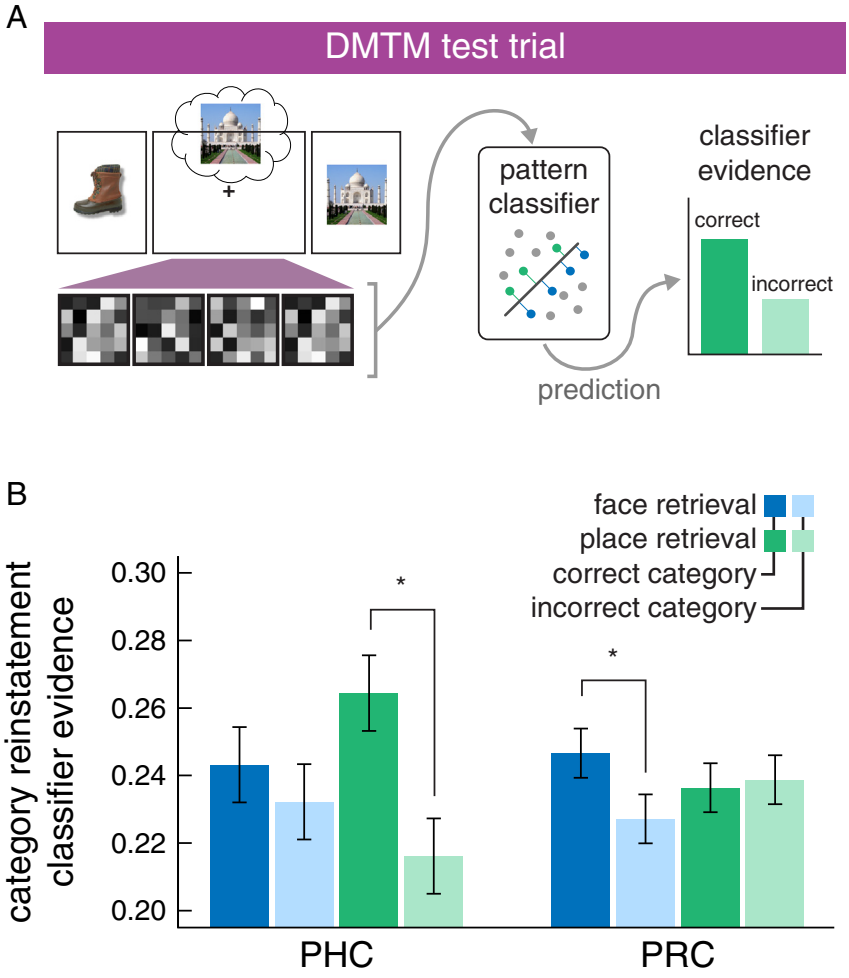

Fig. 3. Schematic of category-level reinstatement decoding analysis and results in PHC and PRC. (A) For each ROI, a linear SVM pattern classifier was first trained on activation patterns acquired during the pre-exposure phase. The classifier was trained to discriminate patterns associated with faces and places and then applied to the data acquired during the delay period of the DMTM trials. Probability estimates for the face and place labels were extracted from the SVM classifier and used to index the amount of classifier evidence for the two categories. Finally, the evidence was relabeled as correct versus incorrect category according to the retrieval category of the trial (in the depicted example, scene would be the correct category). (B) Category-specific reinstatement was observed in MTL cortex, with PRC showing significantly more face than place evidence during face retrieval and PHC showing significantly more place than face evidence during place retrieval. The hippocampus showed no evidence for category reinstatement (see Methods). Error bars represent 95% high density intervals of the paired BEST comparisons between correct and incorrect category classifier evidence. Significant paired differences ($p < 0.05$) are noted with asterisks (*).

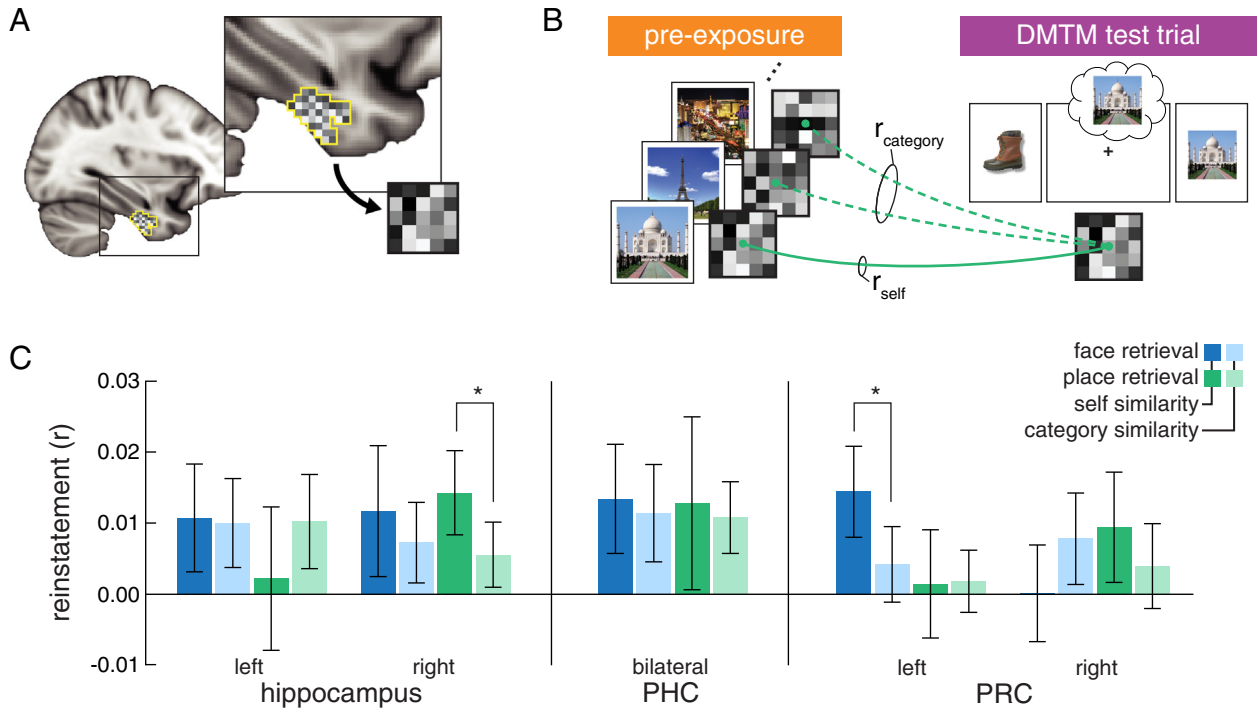

Fig. 4. Schematic of item reinstatement neural pattern similarity analyses and results for each of the ROIs. (A) Within each ROI (perirhinal cortex is depicted in the figure and inset), the pattern of neural activation was extracted for every trial, here represented as a matrix grid. (B) Item-level reinstatement was measured by calculating the similarity between activation patterns during DMTM trials for specific faces and places and activation patterns of viewing the same face or place during the pre-exposure phase (r_{self}) relative to activation patterns of viewing other within-category items ($r_{category}$). (C) Item reinstatement similarity analysis results. For each ROI, self (darker bars) versus category (lighter bars) similarity is plotted for face (blue) and place (green) retrieval trials. Both hippocampus and PRC showed significant differences across hemispheres, thus left and right ROIs within these regions were considered separately. Right hippocampus and left PRC showed item-level reinstatement for place and face retrieval, respectively. Error bars represent 95% confidence intervals of the means. Asterisks (*) note significant paired comparisons of self greater than category similarity corrected for false discovery rate ($\alpha = 0.05$).

that exhibited item-specific reinstatement exclusively for faces or places with a searchlight analysis extension of the item reinstatement analysis (Table 1 and Fig. 5A). Average delay-period activation within these face and place reinstatement regions, as well as right hippocampus and left PRC, were extracted for each correct trial. We estimated the degree of correlation between the MTL ROIs and face/place reinstatement regions separately for correct face and place retrieval trials (Fig. 5B). We found that the connectivity pattern between right hippocampus and face and place reinstatement regions (Fig. 5C) was modulated by the type of retrieval such that there was greater coupling with face regions during face retrieval and place regions during place retrieval ($F_{1,22} = 11.56$, $p = 0.0026$, $\eta_p^2 = 0.344$; see Supplemental Materials for exploratory analysis with right CA_1). In contrast, left PRC connectivity was not modulated by retrieval ($F_{1,22} = 0.038$, $p = 0.85$, $\eta_p^2 = 0.002$).

Table 1
Item reinstatement regions in occipito-temporal cortices. Regions were identified from the paired comparison of face versus place reinstatement with a voxelwise threshold of $p = 0.01$ and cluster extent threshold of $p = 0.05$ (extent greater than 29 voxels). In the table, clusters that survived correction are described by their corresponding region (region), the number of voxels in the cluster (voxels), the t value of the peak voxel in the cluster (t value), and the location of the peak voxel in MNI coordinates (location).

region	voxels	t value	location (x, y, z)
face regions			
left lingual gyrus	87	3.49	−24, −54, −2
right middle temporal	49	2.84	62, −8, −16
left temporal pole	49	3.24	−50, 6, −40
scene regions			
left lateral occipital	727	5.44	−16, −94, 8
right lateral occipital	566	3.77	36, −76, 6
	89	3.6	44, −80, −10
left ventral temporal	143	3.31	−44, −22, −18

Modeling the influence of item reinstatement on decision making

We next sought to assess the behavioral significance of reinstated hippocampal and PRC patterns, specifically testing the hypothesis that reinstated episodic information in hippocampal and PRC patterns serves as the critical evidence for memory-based decision making. We fit participants' choice behavior to the memory probe with the drift diffusion model (DDM). The DDM simultaneously accounts for response choices and times across different experimental conditions by operationally defining decision making as an evidence accumulation process (Ratcliff, 1978). We linked the neural signatures of item reinstatement in right hippocampus and left PRC to the computational mechanisms of decision making in the DDM (Fig. 6A and Table 2).

Table 2
Drift diffusion model (DDM) results. Separate models that included trial-by-trial effects of item reinstatement from the three regions that showed significant item reinstatement in the pattern similarity analysis (right hippocampus and left PRC) on the three core DDM parameters (drift rate v, threshold a, and nondecision time t) were fit and compared to baseline models with no item reinstatement effects (top row). In the table, models are labeled according to their reinstatement-parameter relationship (e.g., the $v \sim r_{hipp}$ model included an effect of hippocampal item reinstatement on drift rate v). Behavior from face (left two columns) and place (right two columns) retrieval conditions were fit in separate models. Rows in bold text note item reinstatement models with a significantly better account of the behavior data than the corresponding baseline models.

Face retrieval trials	DIC	Place retrieval trials	DIC
Baseline model	*−843.09*	*Baseline model*	*−400.12*
$v \sim r_{hipp}$	−842.47	**$v \sim r_{hipp}$**	**−406.32**
$a \sim r_{hipp}$	−844.74	$a \sim r_{hipp}$	−387.15
$t \sim r_{hipp}$	−830.79	$t \sim r_{hipp}$	−399.13
$v \sim r_{PRC}$	**−853.02**	$v \sim r_{PRC}$	−393.14
$a \sim r_{PRC}$	−844.25	$a \sim r_{PRC}$	−400.79
$t \sim r_{PRC}$	−848.36	$t \sim r_{PRC}$	−399.93

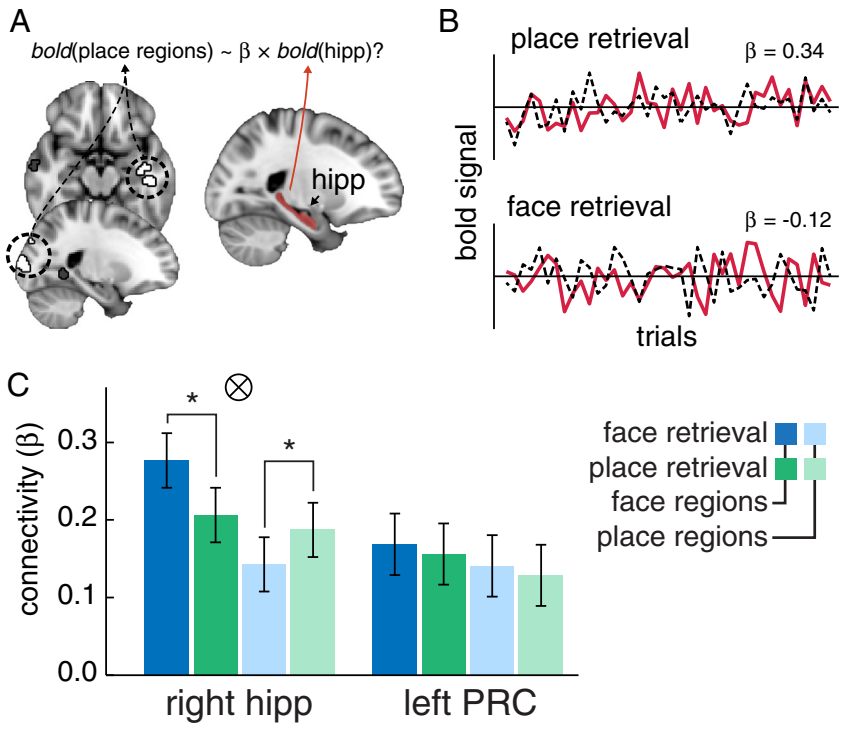

Fig. 5. Schematic of the functional connectivity analysis and the results for right hippocampus and left PRC. (A) Trial-by-trial average activation was extracted from the three MTL ROIs (hippocampus is shown in red on the right) and regions within occipito-temporal cortex (OTC) that exhibited item-reinstatement signatures exclusively for faces or places (face regions shown in dark gray, place regions in white; see Table 1 for a list of all OTC item reinstatement regions). (B) Functional connectivity between the three MTL ROIs and OTC reinstatement regions was assessed with robust regression for face and place retrieval trials separately. In the depicted example, average activation within the right hippocampus (red line) correlates more with average activation in OTC place reinstatement regions (dotted line) during place (top plot) than face retrieval trials (bottom plot). (C) Results for each MTL ROI are plotted in separate graphs that depict average connectivity to OTC regions exhibiting face (darker colors) and place (lighter colors) reinstatement during face (blue) and place (green) retrieval trials. Error bars represent the 95% confidence intervals of the interaction between retrieval type and OTC region. Significant interactions are noted with a tensor symbol, and significant pairwise BEST comparisons with an asterisk.

Specifically, on a trial-by-trial basis, we estimated in separate models the degree that item-specific neural reinstatement predicted evidence accumulation, decision threshold, or non-decision time for match and mismatch probe decisions in the DMTM task using Bayesian parameter estimation (Wiecki et al., 2013). We found that item-specific reinstatement in right hippocampus significantly predicted the rate of evidence accumulation to matching place probes on a trial-by-trial basis (baseline model: $DIC = -400.12$; hippocampus model: $DIC = -406.32$, $P_{match} = 0.021$; Fig. 6B). The relationship between hippocampal reinstatement and evidence accumulation was such that reinstating more place-specific information led to faster evidence accumulation, resulting in faster response times. This relationship was not present for mismatching place probes ($P_{mismatch}$s > 0.4). Item-specific reinstatement in PRC also showed a relationship to evidence accumulation, with greater reinstatement of face-specific information leading to faster accumulation as a general effect across both match and mismatch probes (baseline model: $DIC = -843.09$; PRC model: $DIC = -852.02$, $P = 0.024$). See the Supplemental Materials for additional analyses of the relationship between hippocampal reinstatement and response times (Figure S1), the generative nature of the neurally-informed DDM results (Figure S2), and an exploratory analysis with right CA_1.

We also tested item reinstatement in the OTC regions identified in the searchlight analysis (see Table 1) using the DDM framework. However, none of models that included trial-by-trial reinstatement from OTC place reinstatement regions fit better than the baseline scene model (DICs > −395). One model including reinstatement from a face region in left lingual gyrus as a factor on drift rate did fit slightly better than the face baseline model ($DIC = -845.44$); however, the regression coefficient between item reinstatement and drift rate was not significant (Ps > 0.2). All other models including OTC face reinstatement regions did not fit better than the baseline model (DICs > −843). Although accuracy was near ceiling in the DMTM task, thereby limiting the DDM models to accounting mostly for differences in response time distributions, these results suggest a trial-by-trial link between the fidelity of reinstated memories in the MTL and decisions about the specific contents of those memories.

Discussion

Taken together, our results suggest that the specific contents of episodic memories are reinstated in the activation patterns of hippocampal subfields and subregions of MTL cortex. In particular, our findings that reinstated neural patterns in the hippocampus contain item-level information are consistent with computational theories that point specifically to CA_3 and CA_1 as critical for reinstating complete memory representations from partial sensory input through pattern completion (Marr, 1971; Norman and O'Reilly, 2003; Rolls, 2013).

Importantly, our study extends beyond the existing evidence for hippocampal pattern completion from rodent and human studies that rely on representations of spatial locations (Kyle et al., 2015; Leutgeb et al., 2007; Miller et al., 2013; Neunuebel and Knierim, 2014; Stokes et al., 2015). In particular, such work has been limited in demonstrating the link between hippocampal-based spatial codes and choice behavior. It has been shown that single unit place cell responses can predict direction choices made by rodents in maze tasks (Pastalkova et al., 2008; Singer et al., 2013); however, these findings are not based on pattern completion processes per se. One recent rodent study measured both CA_3 input from DG and CA_3 output, finding direct evidence for pattern completion processes in CA_3, but without a link to behavior (Neunuebel and Knierim, 2014). Finally, none of this research has indexed specific memory content beyond spatial locations (i.e., place fields). By linking a continuous neural measure of memory to a

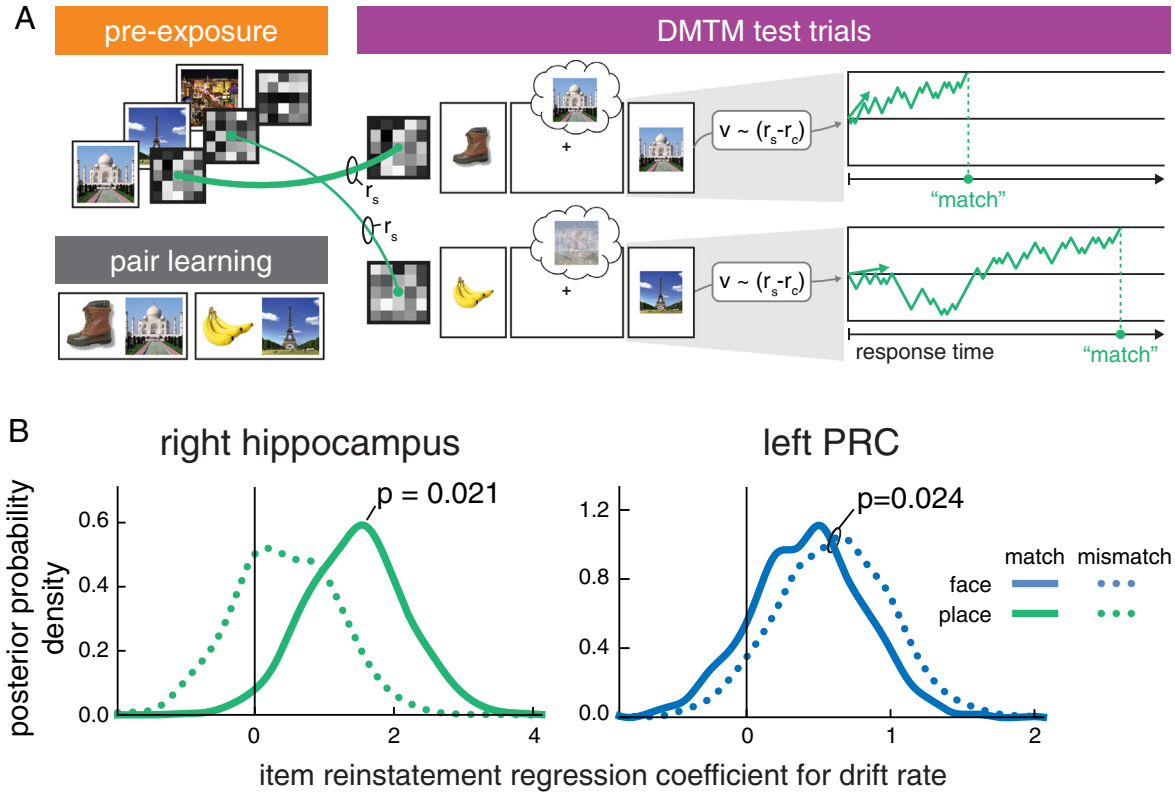

Fig. 6. Schematic of the drift diffusion modeling (DDM) analysis and the results for right hippocampus, right CA_1, and left PRC. (A) An item-level reinstatement index was calculated on a trial-by-trial basis by taking the difference between the self and category similarity on each trial (r_s–r_c). This trial-by-trial item reinstatement index was then entered in a DDM analysis of response choices and times to the probe stimulus. Specifically, the item reinstatement index was modeled as a linear effect on the drift rate (v) parameter of the DDM, thereby influencing the rate of evidence accumulation. For example, the top trial depicts a situation where reinstatement of the place associated with the boot, Taj Mahal, is high. This results in a faster accumulation of evidence and a shorter response time. When item reinstatement is low, as depicted in the bottom trial, evidence accumulates gradually and the response time is longer. (B) The plotted distributions represent the Bayesian posterior densities of the interaction of drift rate and item reinstatement for right hippocampus (left), right CA_1 (middle) and left PRC (right). Higher item reinstatement in right hippocampus and CA_1 increased drift rates on place match trials (solid green), but had no effect on place mismatch trials (dotted green). Item reinstatement in left PRC showed a main effect of facilitation on drift rates for both match and mismatch face trials (blue lines).

prominent mathematical model of decision making, our results provide support for theories of hippocampal function that link retrieval of specific memory elements to reinstatement of hippocampal representations that themselves are critical for subsequent decisions (Lisman and Grace, 2005; Marr, 1971; Norman and O'Reilly, 2003; O'Reilly and Rudy, 2001; Rolls, 2013).

Evidence of memory reinstatement in human hippocampus has been limited to electrophysiological studies of epilepsy patients (e.g., Gelbard-Sagiv et al., 2008; Paz et al., 2010; Rutishauser et al., 2015). By recording single cells within the hippocampus, it has been demonstrated that a subset of cells show item-specific selectivity during visual presentation (Quiroga et al., 2005, 2007), and that selective firing for specific episodes during memory encoding reoccurred during free recall of the same episodes (Gelbard-Sagiv et al., 2008). Furthermore, the strength of episodic encoding in single hippocampal cells, measured as a correlation in temporal patterns of encoding-related firing rates, has been shown to predict subsequent recall performance (Paz et al., 2010). Single cell firing patterns in the hippocampus have also been tied to memory decisions during retrieval, with evidence that a subset of hippocampal cells fire proportionally to graded responses of memory strength and confidence (Rutishauser et al., 2015). This work offers compelling demonstrations of hippocampal memory reinstatement with links to behavior, yet such findings are limited to characterizing the reinstatement signatures of single cells. Moreover, although this work suggest some hippocampal cells represent memory information that is correlated with decision making, the existing findings are limited to a general level of memory information and cannot speak to item-specific pattern completion.

In contrast to electrophysiological studies, fMRI-based findings in support of memory reinstatement in the hippocampus are considerably lacking. One recent fMRI study found evidence for reinstatement of specific memories in hippocampus (Wimber et al., 2015). However, that finding does not address the critical questions set forth in the current study: reinstatement of specific memory contents and their impact on decisions. In particular, the fact that reinstatement in that study was measured by comparing retrieval patterns to episode-specific templates acquired after both encoding and retrieval necessarily reflects retrieval of entire experiences rather than individual event elements. Thus, similar to other studies (Chadwick et al., 2010, 2014), the Wimber et al. (2015) study measured the distinctiveness of individual composite events during retrieval, but not retrieval of the individual memory contents themselves, which is the critical theoretical prediction of pattern completion. More importantly, they failed to link episode-specific hippocampal patterns to behavior, which is the second key theoretical test. In the current study, by comparing reinstated memory representations to visual representations recorded before associative learning, we provide fundamental evidence that hippocampus and perirhinal cortex uniquely reinstate the specific contents of episodic memories. We not only relate this reinstatement to behavior, but also use reinstatement to make trial-by-trial predictions about decisions in a generative fashion.

We observed behaviorally relevant content-specific reinstatement in the hippocampus and PRC. This finding is consistent with proposals that the hippocampus, PRC, and PHC play unique roles in encoding and retrieval depending on the content of memories (e.g., Bird and Burgess, 2008; Diana et al., 2007). For example, the observation that

reinstatement of specific faces in PRC influenced memory decisions converges with prior studies of lesion patients that suggest a selective role for PRC in processing faces (Mundy et al., 2013). Similar research has also implicated the hippocampus in spatial processing with hippocampal lesions leading to deficits in place but not face processing (Bird and Burgess, 2008). Such a functional role for hippocampus is consistent with our findings of item-reinstatement in the hippocampus guiding memory decisions for places.

PHC is also known to be content-selective, with a preference for processing places (Diana et al., 2007). A recent study with a pattern similarity approach similar to the current study found evidence of episode-specific reinstatement of place images in PHC, but not the hippocampus (Staresina et al., 2012). However, PHC is known to code for scene subcategories (Walther et al., 2009) and the restricted set of four place images each representative of disparate place categories (home, office, city, and nature) used in this work suggests these findings could be due to reinstatement of place subcategory rather than episode-specific information. Such an interpretation is consistent with our findings of category, but not item reinstatement in PHC. Although the current study is limited to eight place images, we selected famous landmarks that were not representative of separate place subcategories to better target item rather that category representations.

The relationship between the item-level reinstatement and category-level decoding findings within individual regions also has important implications for the nature of memory representations in each region. Significant category-level reinstatement in PRC and PHC suggests MTL cortical representations are categorically organized such that items from the same category are represented in similar parts of the network, resulting in functional memory modules (Diana et al., 2008; Huffman and Stark, 2014; LaRocque et al., 2013; Liang et al., 2013). This finding stands in clear contrast to the lack of category reinstatement in the hippocampus. The failure of prior decoding approaches to identify hippocampal reinstatement may be due to the relatively category-agnostic nature of the hippocampus such that items from the same category are uniquely represented (Huffman and Stark, 2014; Kuhl and Chun, 2014; LaRocque et al., 2013; Liang et al., 2013).

Hippocampal coding is generally thought to be sparse in nature, a theoretical position with much empirical support (Quiroga et al., 2008; Waydo et al., 2006). Although several fMRI-based studies, in addition to the current study, have successfully measured episodic-specific neural representations in the hippocampus (Bonnici et al., 2012; Chadwick et al., 2010, 2014; Hassabis et al., 2009; Wimber et al., 2015), the sparse representations of the hippocampus are a high methodological hurdle for BOLD-based empirical investigations and may have been a limiting factor in previous attempts to uncover item-specific reinstatement signatures in the hippocampus. An additional factor to consider is that our study included relatively more associative pair training (average of 3.5 presentations of each pair during training across participants) compared to previous studies potentially leading to stronger representations. It may be that more training during encoding is required to observe BOLD-based signatures of hippocampal item reinstatement. However, the extensive associative pair training and the intentional nature of the encoding demands during training is a limitation of the current study; such strong associative representations may have been observable with fMRI methods, but may involve more learning related rather than episodic memory related processes. Nonetheless, the findings in the current study suggest that whereas a hierarchy of category and item-specific information is represented in PRC, hippocampal representations code for the specific contents of memories in a non-modular fashion.

By targeting MTL and OTC regions that exhibit item-specific reinstatement, we were able to uncover a hippocampal-OTC network that demonstrated connectivity patterns that are modulated by current retrieval demands. This finding is consistent with recent rodent studies employing optogenetics to examine functional coupling between hippocampus and behaviorally-relevant cortical regions. In one such study, cortical neurons active during fear conditioning failed to reactivate when CA_1 neurons were silenced during retrieval (Tanaka et al., 2014). Moreover, another study found that CA_1 neurons exhibited distinct projections to different target brain areas depending on the type of task performed by the animal (Ciocchi et al., 2015). One recent human fMRI study examined the connectivity between the hippocampus and MTL cortex with dynamic causal modeling, finding that hippocampal activation mediates successful retrieval of associative information from content-specific regions of MTL (Staresina et al., 2013). The hippocampus has also been implicated in several fMRI studies examining cortical reinstatement. Hippocampal activation during both encoding and retrieval scales with task (Leiker and Johnson, 2015), category (Gordon et al., 2014; Horner et al., 2015) and item (Ritchey et al., 2012; Wing et al., 2015) reinstatement signatures in neocortex. In particular, Ritchey et al. (2012) found a relationship between hippocampal univariate activation and the degree of encoding-retrieval pattern similarity in cortical regions, consistent with the notion that hippocampus mediates memory reinstatement in cortex. However, because the same visual input was present at both encoding and retrieval, it is unknown whether their findings are related to reinstatement processes or to perceptual processes acting on the same visual input during both encoding and retrieval.

The current findings build upon and extend these previous results by demonstrating that during retrieval, hippocampus is preferentially coupled with OTC regions that reinstate item-level representations of memory content. It should be noted that the current connectivity findings are limited by their correlative nature; specifically, the task-dependent correlation between hippocampal and OTC neural activity cannot be used to infer the existence or directionality of a causal relationship between the brain regions (Friston, 2011). However, the current findings, together with the existing literature, are consistent with the theoretical proposal that the hippocampus plays an important role in the re-experiencing of mnemonic content through cortical reinstatement and may route behaviorally-relevant information to task-specific cortical areas during memory retrieval.

It is worth noting that although the current work has focused on reinstatement as arising from pattern completion, pattern separation is also functionally tied to memory reinstatement. In particular, overlapping experiences that share content must be encoded in orthogonal representations in order to retrieve distinct memory traces of these experiences (Marr, 1971; Norman and O'Reilly, 2003; O'Reilly and Rudy, 2001). This separation process is driven by competition between incoming sensory information from the current experience and reinstated information from related previously encoded memories (Hulbert and Norman, 2014; Norman and O'Reilly, 2003; O'Reilly and Rudy, 2001). The current study was designed to target reinstatement signatures arising from pattern completion processes. The object cues during the DMTM task provided partial information for retrieving the correctly paired face or place associate through pattern completion. It is likely that pattern separation was involved both during encoding of the paired associates and potentially influenced what information was available during retrieval (Kyle et al., 2015). However, the specificity of the item reinstatement signatures we observed during retrieval is consistent with a "filling in" of item-specific mnemonic information through pattern completion. Future studies could extend these findings by parametrically varying the degradation or similarity between a cue and the associative memory (e.g., Bakker et al., 2008; Kirwan and Stark, 2007) to characterize how degraded cues lead to reinstatement of item specific information through pattern completion.

The present findings provide a critical demonstration for the theorized role of reinstated hippocampal representations in mnemonic decision making (Lisman and Grace, 2005; Norman and O'Reilly, 2003), while further revealing important information about the organization of the underlying mnemonic codes. Theoretical accounts posit that

reinstated representations in CA_3 and CA_1 act as predictions about present events derived from past knowledge. Limited support for this proposal is found in existing neuroimaging work that suggests the magnitude of hippocampal activation during retrieval is correlated with the speed of memory decisions (Gordon et al., 2014) and confidence judgments (Leiker and Johnson, 2015; Thakral et al., 2015). In the current study, we provided a direct test of the role of the hippocampus in memory retrieval by linking trial-by-trial item reinstatement signatures to the accumulation of evidence in a computational model of memory-based decisions. Indeed, the computational modeling analysis showed that the fidelity of reinstated item-level representations is directly predictive of the speed of memory decisions. The relationship between reinstatement and behavior was specific to MTL-based item reinstatement; reinstatement in occipito-temporal regions was not predictive of mnemonic decisions. Mechanistically, these findings are consistent with a memory-based decision process whereby the memory evidence internally generated during retrieval in the hippocampus and perirhinal cortex is compared to the sensory input of the probe stimulus. The fidelity of the memory evidence influences this comparison such that more evidence leads to fast accumulation of decision evidence towards a match/mismatch response.

It is interesting to note that although the DDM has been leveraged more recently as a model of perceptual decision making (Ratcliff and Rouder, 1998), its original conception was an account of memory retrieval that linked internally-generated information to the accumulation of evidence for memory-based decisions (Ratcliff, 1978). Our results bridge brain measures and behavior to ground this classic theory's mechanisms in neural evidence and show that internally generated content from past experiences guides decision making. Moreover, our approach opens the door to answering long-standing questions about memory function at the resolution of individual memories in healthy individuals. For instance, how do specific memories emerge over learning experiences? When are the details of the past forgotten versus successfully consolidated into robust memory networks? How are the specific elements of memories flexibly integrated to create new knowledge? The current study also paves the way towards a unified computational approach in which the neural representations supporting our memories are quantitatively linked to the formal mechanisms that underlie decision making.

Acknowledgements

The authors declare no competing financial interests. This study was supported by NIH (F32-MH100904), NSF CAREER Award (1056019), and NIH (R01-MH100121). We wish to thank A. Huk for helpful comments on an earlier version of this report; J. Mumford for advice on optimizing the experimental design for RSA; B. Love for helpful comments on the DDM analysis; and M. Schlichting, E. Stein, B. Gelman, and K. Nguyen for their help with data collection. Thank you also to the Texas Advanced Computing Center (TACC; http://www.tacc.utexas.edu) at The University of Texas at Austin for providing high-performance computing resources critical for the work presented within this report.

Appendix A. Supplementary data

Supplementary data to this article can be found online at http://dx.doi.org/10.1016/j.neuroimage.2015.12.015.

References

Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C., 2011. A reproducible evaluation of ANTs similarity metric performance in brain image registration. NeuroImage 54, 2033–2044. http://dx.doi.org/10.1016/j.neuroimage.2010.09.025.

Bakker, A., Kirwan, C.B., Miller, M., Stark, C.E.L., 2008. Pattern separation in the human hippocampal CA3 and dentate gyrus. Science 319, 1640–1642. http://dx.doi.org/10.1126/science.1152882.

Bird, C.M., Burgess, N., 2008. The hippocampus and memory: insights from spatial processing. Nat. Rev. Neurosci. 9, 182–194. http://dx.doi.org/10.1038/nrn2335.

Bogacz, R., Wagenmakers, E.J., Forstmann, B.U., Nieuwenhuis, S., 2010. The neural basis of the speed-accuracy tradeoff. Trends Neurosci. 33, 10–16. http://dx.doi.org/10.1016/j.tins.2009.09.002.

Bollimunta, A., Totten, D., Ditterich, J., 2012. Neural dynamics of choice: single-trial analysis of decision-related activity in parietal cortex. J. Neurosci. 32, 12684–12701. http://dx.doi.org/10.1523/JNEUROSCI.5752-11.2012.

Bonnici, H.M., Chadwick, M.J., Kumaran, D., Hassabis, D., Weiskopf, N., Maguire, E.A, 2012. Multi-voxel pattern analysis in human hippocampal subfields. Front. Hum. Neurosci. 6, 1–13. http://dx.doi.org/10.3389/fnhum.2012.00290.

Bosch, S.E., Jehee, J.F.M., Fernández, G., Doeller, C.F., 2014. Reinstatement of associative memories in early visual cortex is signaled by the hippocampus. J. Neurosci. 34, 7493–7500. http://dx.doi.org/10.1523/JNEUROSCI.0805-14.2014.

Chadwick, M.J., Hassabis, D., Weiskopf, N., Maguire, E.A., 2010. Decoding individual episodic memory traces in the human hippocampus. Curr. Biol. 1–4.

Chadwick, M.J., Hassabis, D., Maguire, E.A., 2011. Decoding overlapping memories in the medial temporal lobes using high-resolution fMRI. Learn. Mem. 8, 742–746. http://dx.doi.org/10.1101/lm.023671.111.

Chadwick, M.J., Bonnici, H.M., Maguire, E.A., 2014. CA3 size predicts the precision of memory recall. Proc. Natl. Acad. Sci. 111, 10720–10725. http://dx.doi.org/10.1073/pnas.1319641111.

Chen, J., Olsen, R.K., Preston, A.R., Glover, G.H., Wagner, A.D., 2011. Associative retrieval processes in the human medial temporal lobe: hippocampal retrieval success and CA1 mismatch detection. Learn. Mem. 18, 523–528. http://dx.doi.org/10.1101/lm.2135211.

Ciocchi, S., Passecker, J., Malagon-Vina, H., Mikus, N., Klausberger, T., 2015. Selective information routing by ventral hippocampal CA1 projection neurons. Science 348, 560–563.

Diana, R.A., Yonelinas, A.P., Ranganath, C., 2007. Imaging recollection and familiarity in the medial temporal lobe: a three-component model. Trends Cogn. Sci. 11, 379–386. http://dx.doi.org/10.1016/j.tics.2007.08.001.

Diana, R.A., Yonelinas, A.P., Ranganath, C., 2008. High-resolution multi-voxel pattern analysis of category selectivity in the medial temporal lobes. Hippocampus 18, 536–541. http://dx.doi.org/10.1002/hipo.20433.

Eichenbaum, H., 2004. Hippocampus: cognitive processes and neural representations that underlie declarative memory. Neuron 44, 109–120. http://dx.doi.org/10.1016/j.neuron.2004.08.028.

Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., Van Der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., Dale, A.M., 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. Neuron 33, 341–355. http://dx.doi.org/10.1016/S0896-6273(02)00569-X.

Friston, K.J., 2011. Functional and effective connectivity: a review. Brain Connect. 1, 13–36. http://dx.doi.org/10.1089/brain.2011.0008.

Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., Fried, I., 2008. Internally generated reactivation of single neurons in human hippocampus during free recall. Science 322, 96–101. http://dx.doi.org/10.1126/science.1164685.

Glosser, G., Saykin, A.J., Deutsch, G.K., O'Connor, M.J., 1995. Neural organization of material-specific memory functions in temporal lobe epilepsy patients as assessed by the intracarotid amobarbital test. Neuropsychology 9, 449–456. http://dx.doi.org/10.1037/0894-4105.9.4.449.

Golby, A.J., Poldrack, R.A., Brewer, J.B., Spencer, D., Desmond, J.E., Aron, A.P., Gabrieli, J.D., 2001. Material-specific lateralization in the medial temporal lobe and prefrontal cortex during memory encoding. Brain 124, 1841–1854. http://dx.doi.org/10.1093/brain/124.9.1841.

Gold, J.I., Shadlen, M.N., 2007. The neural basis of decision making. Annu. Rev. Neurosci. 30, 535–574. http://dx.doi.org/10.1146/annurev.neuro.29.051605.113038.

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. Cereb. Cortex 24, 3350–3364. http://dx.doi.org/10.1093/cercor/bht194.

Hanke, M., Halchenko, Y.O., Sederberg, P.B., Hanson, S.J., Haxby, J.V., Pollmann, S., 2009. PyMVPA: a python toolbox for multivariate pattern analysis of fMRI data. Neuroinformatics 7, 37–53. http://dx.doi.org/10.1007/s12021-008-9041-y.

Hanks, T.D., Kiani, R., Shadlen, M.N., 2014. A neural mechanism of speed-accuracy tradeoff in macaque area LIP. Elife 2014, 1–17. http://dx.doi.org/10.7554/eLife.02260.

Hassabis, D., Chu, C., Rees, G., Weiskopf, N., Molyneux, P.D., Maguire, E.A., 2009. Decoding neuronal ensembles in the human hippocampus. Curr. Biol. 19, 546–554. http://dx.doi.org/10.1016/j.cub.2009.02.033.

Heekeren, H.R., Marrett, S., Ungerleider, L.G., 2008. The neural systems that mediate human perceptual decision making. Nat. Rev. Neurosci. 9, 467–479. http://dx.doi.org/10.1038/nrn2374.

Heit, G., Smith, M.E., Halgren, E., 1988. Neural encoding of individual words and faces by the human hippocampus and amygdala. Nature 333, 773–775. http://dx.doi.org/10.1038/333773a0.

Horner, A.J., Bisby, J.A., Bush, D., Lin, W.-J., Burgess, N., 2015. Evidence for holistic episodic recollection via hippocampal pattern completion. Nat. Commun. 6, 7462. http://dx.doi.org/10.1038/ncomms8462.

Huffman, D.J., Stark, C.E.L., 2014. Multivariate pattern analysis of the human medial temporal lobe revealed representationally categorical cortex and representationally agnostic hippocampus. Hippocampus 10, 1–10. http://dx.doi.org/10.1002/hipo.22321.

Hulbert, J.C., Norman, K.A., 2014. Neural differentiation tracks improved recall of competing memories following interleaved study and retrieval practice. Cereb. Cortex 1–15. http://dx.doi.org/10.1093/cercor/bhu284.

Johnson, J.D., McDuff, S.G.R., Rugg, M.D., Norman, K.A., 2009. Recollection, familiarity, and cortical reinstatement: a multivoxel pattern analysis. Neuron 63, 697–708. http://dx.doi.org/10.1016/j.neuron.2009.08.011.

Kamitani, Y., Sawahata, Y., 2010. Spatial smoothing hurts localization but not information: pitfalls for brain mappers. NeuroImage 49, 1949–1952. http://dx.doi.org/10.1016/j.neuroimage.2009.06.040.

Kirwan, C.B., Stark, C.E.L., 2007. Overcoming interference: an fMRI investigation of pattern separation in the medial temporal lobe. Learn. Mem. 14, 625–633. http://dx.doi.org/10.1101/lm.663507.

Knierim, J.J., Neunuebel, J.P., Deshmukh, S.S., B, P.T.R.S., 2014. Cortex : Objects, Path Integration And Local—Global Reference Frames Functional Correlates Of The Lateral And Medial Entorhinal Cortex : Objects, Path Integration And Local–Global Reference Frames.

Kriegeskorte, N., Mur, M., Bandettini, P., 2008. Representational similarity analysis — connecting the branches of systems neuroscience. Front. Syst. Neurosci. 2, 4. http://dx.doi.org/10.3389/neuro.06.004.2008.

Kruschke, J.K., 2012. Bayesian estimation supersedes the t test. J. Exp. Psychol. Gen. 142, 573–603. http://dx.doi.org/10.1037/a0029146.

Kuhl, B.A., Chun, M.M., 2014. Successful remembering elicits event-specific activity patterns in lateral parietal cortex. J. Neurosci. 34, 8051–8060. http://dx.doi.org/10.1523/JNEUROSCI.4328-13.2014.

Kyle, C.T., Stokes, J.D., Lieberman, J.S., Hassan, A.S., Ekstrom, A.D., 2015. Successful retrieval of competing spatial environments in humans involves hippocampal pattern separation mechanisms. Elife 4, 160. http://dx.doi.org/10.7554/eLife.10499.

Lacy, J.W., Yassa, M.A., Stark, S.M., Muftuler, L.T., Stark, C.E.L., 2011. Distinct pattern separation related transfer functions in human CA3/dentate and CA1 revealed using high-resolution fMRI and variable mnemonic similarity. Learn. Mem. 18, 15–18. http://dx.doi.org/10.1101/lm.1971111.

LaRocque, K.F., Smith, M.E., Carr, V.a., Witthoft, N., Grill-Spector, K., Wagner, A.D., 2013. Global similarity and pattern separation in the human medial temporal lobe predict subsequent memory. J. Neurosci. 33, 5466–5474. http://dx.doi.org/10.1523/JNEUROSCI.4293-12.2013.

Leiker, E.K., Johnson, J.D., 2015. Pattern reactivation CO-varies with activity in the core recollection network during source memory. Neuropsychologia 75, 88–98. http://dx.doi.org/10.1016/j.neuropsychologia.2015.05.021.

Leutgeb, J.K., Leutgeb, S., Moser, M.-B., Moser, E.I., 2007. Pattern separation in the dentate gyrus and CA3 of the hippocampus. Science 315, 961–966. http://dx.doi.org/10.1126/science.1135801.

Liang, J.C., Wagner, A.D., Preston, A.R., 2013. Content representation in the human medial temporal lobe. Cereb. Cortex 23, 80–96. http://dx.doi.org/10.1093/cercor/bhr379.

Lisman, J.E., Grace, A.A., 2005. The hippocampal-VTA loop: controlling the entry of information into long-term memory. Neuron 46, 703–713. http://dx.doi.org/10.1016/j.neuron.2005.05.002.

Marr, D., 1971. Simple memory: a theory for archicortex. Philos. Trans. R. Soc. Lond. Ser. B Biol. Sci. 262, 23–81. http://dx.doi.org/10.1098/rstb.1971.0078.

McClelland, J.L., McNaughton, B.L., 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. http://dx.doi.org/10.1037/0033-295X.102.3.419.

Miller, J.F., Neufang, M., Solway, A., Brandt, A., Trippel, M., Mader, I., Hefft, S., Merkow, M., Polyn, S.M., Jacobs, J., Kahana, M.J., Schulze-Bonhage, A., 2013. Neural activity in human hippocampal formation reveals the spatial context of retrieved memories. Science 342, 1111–1114. http://dx.doi.org/10.1126/science.1244056.

Mumford, J.A., Turner, B.O., Ashby, F.G., Poldrack, R.A., 2012. Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. NeuroImage 59, 2636–2643. http://dx.doi.org/10.1016/j.neuroimage.2011.08.076.

Mundy, M.E., Downing, P.E., Dwyer, D.M., Honey, R.C., Graham, K.S., 2013. A critical role for the hippocampus and perirhinal cortex in perceptual learning of scenes and faces: complementary findings from amnesia and FMRI. J. Neurosci. 33, 10490–10502. http://dx.doi.org/10.1523/JNEUROSCI.2958-12.2013.

Neunuebel, J.P., Knierim, J.J., 2014. CA3 retrieves coherent representations from degraded input: direct evidence for CA3 pattern completion and dentate gyrus pattern separation. Neuron 81, 416–427. http://dx.doi.org/10.1016/j.neuron.2013.11.017.

Ng, A.Y., Jordan, M.I., 2002. On discriminative vs. Generative classifiers: a comparison of logistic regression and naive bayes, neural information processing systems. Adv. Neural Inf. Process. Syst. 14, 841–849.

Nienborg, H., Cumming, B.G., 2009. Decision-related activity in sensory neurons reflects more than a neuron's causal effect. Nature 459, 89–92.

Norman, K., O'Reilly, R., 2003. Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. Psychol. Rev. 110, 611–646.

Op de Beeck, H.P., 2010. Against hyperacuity in brain reading: spatial smoothing does not hurt multivariate fMRI analyses? NeuroImage 49, 1943–1948. http://dx.doi.org/10.1016/j.neuroimage.2009.02.047.

O'Reilly, R.C., Rudy, J.W., 2001. Conjunctive representations in learning and memory: principles of cortical and hippocampal function. Psychol. Rev. 108, 311–345.

Pastalkova, E., Itskov, V., Amarasingham, A., Buzsáki, G., 2008. Internally generated cell assembly sequences in the rat hippocampus. Science 321, 1322–1327. http://dx.doi.org/10.1126/science.1159775.

Paz, R., Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., Fried, I., 2010. A neural substrate in the human hippocampus for linking successive events. Proc. Natl. Acad. Sci. U. S. A. 107, 6046–6051. http://dx.doi.org/10.1073/pnas.0910834107.

Polyn, S.M., Natu, V.S., Cohen, J.D., Norman, K.A., 2005. Category-specific cortical activity precedes retrieval during memory search. Science 310, 1963–1966. http://dx.doi.org/10.1126/science.1117645.

Quiroga, R.Q., Reddy, L., Kreiman, G., Koch, C., Fried, I., 2005. Invariant visual representation by single neurons in the human brain. Nature 435, 1102–1107. http://dx.doi.org/10.1038/nature03687.

Quiroga, R.Q., Reddy, L., Koch, C., Fried, I., 2007. Decoding visual inputs from multiple neurons in the human temporal lobe. J. Neurophysiol. 98, 1997–2007. http://dx.doi.org/10.1152/jn.00125.2007.

Quiroga, R.Q., Kreiman, G., Koch, C., Fried, I., 2008. Sparse but not "grandmother-cell" coding in the medial temporal lobe. Trends Cogn. Sci. 12, 87–91. http://dx.doi.org/10.1016/j.tics.2007.12.003.

Ratcliff, R., 1978. A theory of memory retrieval. Psychol. Rev. 85, 59–108. http://dx.doi.org/10.1037/0033-295X.85.2.59.

Ratcliff, R., Rouder, J.N., 1998. Modeling response times for two-choice decisions. Psychol. Sci. 9, 347–356.

Rissman, J., Gazzaley, A., D'Esposito, M., 2004. Measuring functional connectivity during distinct stages of a cognitive task. NeuroImage 23, 752–763.

Ritchey, M., Wing, E.A., LaBar, K.S., Cabeza, R., 2012. Neural similarity between encoding and retrieval is related to memory via hippocampal interactions. Cereb. Cortex 23, 2818–2828. http://dx.doi.org/10.1093/cercor/bhs258.

Rolls, E.T., 2013. The mechanisms for pattern completion and pattern separation in the hippocampus. Front. Syst. Neurosci. 7, 74. http://dx.doi.org/10.3389/fnsys.2013.00074.

Rutishauser, U., Ye, S., Koroma, M., Tudusciuc, O., Ross, I.B., Chung, J.M., Mamelak, A.N., 2015. Representation of retrieval confidence by single neurons in the human medial temporal lobe. Nat. Neurosci. 1–12. http://dx.doi.org/10.1038/nn.4041.

Schapiro, A.C., Kustner, L.V., Turk-Browne, N.B., 2012. Shaping of object representations in the human medial temporal lobe based on temporal regularities. Curr. Biol. 22, 1622–1627. http://dx.doi.org/10.1016/j.cub.2012.06.056.

Schlichting, M.L., Preston, A.R., 2014. Memory reactivation during rest supports upcoming learning of related content. Proc. Natl. Acad. Sci. 111, 15845–15850. http://dx.doi.org/10.1073/pnas.1404396111.

Schlichting, M.L., Mumford, J.A., Preston, A.R., 2015. Learning-related representational changes reveal dissociable integration and separation signatures in the hippocampus and prefrontal cortex. Nat. Commun. 6, 8151. http://dx.doi.org/10.1038/ncomms9151.

Singer, A.C., Carr, M.F., Karlsson, M.P., Frank, L.M., 2013. Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task. Neuron 77, 1163–1173. http://dx.doi.org/10.1016/j.neuron.2013.01.027.

Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., Matthews, P.M., 2004. Advances in functional and structural MR image analysis and implementation as FSL. NeuroImage 23 (Suppl. 1), S208–S219. http://dx.doi.org/10.1016/j.neuroimage.2004.07.051.

Squire, L.R., 2004. Memory systems of the brain: a brief history and current perspective. Neurobiol. Learn. Mem. 82, 171–177. http://dx.doi.org/10.1016/j.nlm.2004.06.005.

Staresina, B.P., Henson, R.N.A., Kriegeskorte, N., Alink, A., 2012. Episodic reinstatement in the medial temporal lobe. J. Neurosci. 32, 18150–18156. http://dx.doi.org/10.1523/JNEUROSCI.4156-12.2012.

Staresina, B.P., Cooper, E., Henson, R.N., 2013. Reversible information flow across the medial temporal lobe: the hippocampus links cortical modules during memory retrieval. J. Neurosci. 33, 14184–14192. http://dx.doi.org/10.1523/JNEUROSCI.1987-13.2013.

Stokes, J., Kyle, C., Ekstrom, A.D., 2015. Complementary roles of human hippocampal subfields in differentiation and integration of spatial context. J. Cogn. Neurosci. 27, 546–559. http://dx.doi.org/10.1162/jocn_a_00736.

Tanaka, K.Z., Pevzner, A., Hamidi, A.B., Nakazawa, Y., Graham, J., Wiltgen, B.J., 2014. Cortical representations are reinstated by the hippocampus during memory retrieval. Neuron 84, 347–354. http://dx.doi.org/10.1016/j.neuron.2014.09.037.

Thakral, P.P., Wang, T.H., Rugg, M.D., 2015. Cortical reinstatement and the confidence and accuracy of source memory. NeuroImage 109, 118–129. http://dx.doi.org/10.1016/j.neuroimage.2015.01.003.

Walther, D.B., Caddigan, E., Fei-Fei, L., Beck, D.M., 2009. Natural scene categories revealed in distributed patterns of activity in the human brain. J. Neurosci. 29, 10573–10581. http://dx.doi.org/10.1523/JNEUROSCI.0559-09.2009.

Waydo, S., Kraskov, A., Quian Quiroga, R., Fried, I., Koch, C., 2006. Sparse representation in the human medial temporal lobe. J. Neurosci. 26, 10232–10234. http://dx.doi.org/10.1523/JNEUROSCI.2101-06.2006.

Wiecki, T.V., Sofer, I., Frank, M.J., 2013. HDDM: hierarchical bayesian estimation of the drift-diffusion model in python. Front. Neuroinform. 7, 14. http://dx.doi.org/10.3389/fninf.2013.00014.

Wimber, M., Alink, A., Charest, I., Kriegeskorte, N., Anderson, M.C., 2015. Retrieval induces adaptive forgetting of competing memories via cortical pattern suppression. Nat. Publ. Group 18, 582–589. http://dx.doi.org/10.1038/nn.3973.

Wing, E.A., Ritchey, M., Cabeza, R., 2015. Reinstatement of individual past events revealed by the similarity of distributed activation patterns during encoding and retrieval. J. Cogn. Neurosci. 27, 679–691. http://dx.doi.org/10.1162/jocn_a_00740.

Winkler, A.M., Ridgway, G.R., Webster, M.a., Smith, S.M., Nichols, T.E., 2014. Permutation inference for the general linear model. NeuroImage 92, 381–397. http://dx.doi.org/10.1016/j.neuroimage.2014.01.060.

Xue, G., Dong, Q., Chen, C., Lu, Z., Mumford, J.A., Poldrack, R.A., 2010. Greater neural pattern similarity across repetitions is associated with better memory. Science 330, 97–101. http://dx.doi.org/10.1126/science.1193125.

Zeithamova, D., Dominick, A.L., Preston, A.R., 2012. Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. Neuron 75, 168–179. http://dx.doi.org/10.1016/j.neuron.2012.05.010.