NeuroImage 155 (2017) 513–529

Contents lists available at ScienceDirect

NeuroImage

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

The influence of low-level stimulus features on the representation of contexts, items, and their mnemonic associations

Derek J. Huffman, Craig E.L. Stark*

Department of Neurobiology and Behavior, Center for the Neurobiology of Learning and Memory, University of California, Irvine, United States

ARTICLE INFO

Keywords:
Associative memory
confound
Medial temporal lobe (MTL)
Retrosplenial cortex
Fine-grained semantic representations
Perirhinal cortex (PRC)

ABSTRACT

Since the earliest attempts to characterize the "receptive fields" of neurons, a central aim of many neuroscience experiments is to elucidate the information that is represented in various regions of the brain. Recent studies suggest that, in the service of memory, information is represented in the medial temporal lobe in a conjunctive or associative form with the contextual aspects of the experience being the primary factor or highest level of the conjunctive hierarchy. A critical question is whether the information that has been observed in these studies reflects notions such as a cognitive representation of context or whether the information reflects the low-level sensory differences between stimuli. We performed two functional magnetic resonance imaging experiments to address this question and we found that associative representations observed between context and item (and order) in the human brain can be highly influenced by low-level sensory differences between stimuli. Our results place clear constraints on the experimental design of studies that aim to investigate the representation of contexts and items during performance of associative memory tasks. Moreover, our results raise interesting theoretical questions regarding the disambiguation of memory-related representations from processing-related representations.

Introduction

Following the discovery that removal of structures within the human medial temporal lobe (MTL) causes amnesia (Scoville and Milner, 1957), decades of research have focused on elucidating the contributions of subregions of the MTL to declarative memory. While there is still debate over the precise nature of the division of labor within the MTL, there is consensus that the MTL sits at the apex of a cortical circuit, which allows it to bind the constituents of an event (e.g., "what", "where", "when") into an associative, conjunctive, or relational representation (e.g., "what-what", "what-where"; Mishkin et al., 1997; Cohen et al., 1999; Lavenex and Amaral, 2000; Davachi, 2006; Morris, 2006; Diana et al., 2007; Eichenbaum et al., 2007; Wixted and Squire, 2011; Ranganath and Ritchey, 2012, McKenzie et al., 2015). These theories differ in their emphasis of the role of the hippocampus versus adjacent MTL cortical regions in the formation of such representations—i.e., some theories suggest a more exclusive role for the hippocampus with a more domain-specific involvement of MTL cortical regions—and many studies have begun to test these competing hypotheses.

In a series of groundbreaking studies, Eichenbaum and colleagues used a context-guided object association task to explore how the components of an associative memory such as context, item, position, and valance are represented neurally (Rajji et al., 2006; Komorowski et al., 2009, 2013; Navawongse and Eichenbaum, 2013; Tort et al., 2013; McKenzie et al., 2014; Farovik et al., 2015; Keene et al., 2016). Briefly, animals learn item-reward associations that differ based on the context, which was operationally defined as visually, tactilely, and spatially (side of the apparatus) distinct chambers. Impaired context-guided object association learning has been shown in rats with hippocampal lesions (Komorowski et al., 2013) and in mice with impaired NMDA receptor function in the CA3 subregion of the hippocampus (Rajji et al., 2006), thus establishing a necessary role for the hippocampus in task performance. Recent studies used electrophysiology and representational similarity analysis to investigate patterns of activity across ensembles of cells in the hippocampus (McKenzie et al., 2014) and in MTL cortical regions (Keene et al., 2016). The results of these studies have suggested that subregions of the MTL—including the hippocampus—carry conjunctive representations of the features that comprise an event, including context, item, position, and valence. Moreover, their results suggest that context plays a dominant, organizing role for representations in the MTL, sitting at the highest level of a hierarchy of information (for review see: McKenzie et al., 2015).

* Correspondence to: Department of Neurobiology and Behavior, Center for the Neurobiology of Learning and Memory, University of California, Irvine, 320 Qureshey Research Laboratory, Irvine, CA, United States.
E-mail address: cestark@uci.edu (C.E.L. Stark).

http://dx.doi.org/10.1016/j.neuroimage.2017.04.019
Received 27 October 2016; Accepted 7 April 2017
Available online 08 April 2017

Context is a broad term that encompasses many aspects of an animal's state, including spatial context, temporal context, expectations, and internal state (cf. Nadel and Willner, 1980). Context is an important aspect of memory-driven behavior as it is frequently the case that in one context a certain set of behaviors are appropriate while in another context a different set of behaviors are more appropriate. In many laboratory experiments, context is modeled as a general occasion setter such that the desired behavior or decision to be made is contingent upon or altered if the screen background differs, if the color of the room differs, if the previous time through the corridor you went left (versus right), etc. In the context-guided object association task, the contexts and items are composed of distinct elements (i.e., different visual, olfactory, and tactile cues). Thus a critical issue to address is whether the neural representations of contexts and items reflect the cognitive representation of contexts and items or the low-level sensory differences between stimuli. The use of distinct contexts and items allows animals (e.g., rats, humans) to rapidly discriminate between them (cf. Bulkin et al., 2016); however, as we discovered in the course of the present report, it is virtually inevitable that representational differences will also be present in the relevant primary sensory regions. Given that patterns of activity in the hippocampus (McKenzie et al., 2014) and MTL cortex (Keene et al., 2016) were very dissimilar in response to events that took place in different contexts, we examined whether the representation of context maintains in the absence of low-level sensory differences between contexts. We propose that the cognitive representation of a context should be stable across different versions of the same context (e.g., different viewpoints) so long as the context signals a reliable behavioral outcome (e.g., Context A + Item X + Response 1=Reward).

We developed two human versions of the context-guided object association task for functional magnetic resonance imaging (fMRI) to investigate the representation of context, items, order, and their conjunctions within subregions of the MTL, including the hippocampus, parahippocampal cortex (PHC), and perirhinal cortex (PRC). Additionally, we investigated representations in retrosplenial cortex (RSC), a subregion of the posterior cingulate cortex (PCC), which has been hypothesized to be involved in processing scenes and contexts (Chen et al., 1994; Ennaceur et al., 1997; Cho and Sharp, 2001; Vann and Aggleton, 2002; Bar and Aminoff, 2003; Parron and Save, 2004; Park and Chun, 2009; Walther et al., 2009; Auger and Maguire, 2013; Alexander and Nitz, 2015; Auger et al., 2015; Wing et al., 2015) in addition to playing a role in declarative memory, spatial memory, and the formation of stimulus-stimulus associations (Valenstein et al., 1987; Vann et al., 2009; Aggleton, 2010; Ranganath and Ritchey, 2012; Bucci and Robinson, 2014). In Experiment 1, we used distinct stimuli for our contexts and objects, similar to the rodent studies (Rajji et al., 2006; Komorowski et al., 2009, 2013; Navawongse and Eichenbaum, 2013; Tort et al., 2013; McKenzie et al., 2014; Farovik et al., 2015; Keene et al., 2016). In Experiment 2, we matched the low-level visual features of our stimulus set to test for context and object representation in the absence of the confounding effect of low-level sensory differences between stimuli.

The results of Experiment 1 are consistent with the representation of context in the MTL. Additionally, the results of Experiment 1 are consistent with the notion that RSC/PCC carries context and conjunctive item-in-context information and such representations correlated with behavioral performance (a traditional means of enhancing our confidence that the observed signals are mnemonic in nature). However, the results of Experiment 1 are also consistent with the representation of context, item-in-context, and item-in-order-in-context in primary visual cortex. Further, the relationship between such representations and behavior was at least as strong in primary visual cortex as in RSC/PCC. Therefore, our results provide a clear demonstration of the importance of controlling for low-level feature differences between contexts and objects. Moreover, these results raise interesting questions about how to distinguish between memory-related representations and processing-related representations. In Experiment 2, we matched the low-level features between our contexts and objects, and we found that the evidence for context and associative representations disappeared, suggesting that the results from Experiment 1 were influenced by differences in the low-level features that comprised the events. In contrast, we observed evidence for fine-grained object representation in PRC in the absence of a low-level confound, thus corroborating theories that suggest that PRC contains fine-grained semantic representations of objects (e.g., Clarke and Tyler, 2015).

Materials and methods

Participants

Thirty-five participants were recruited from the community at the University of California, Irvine. Participants were between 18 and 31 years of age, were right handed, and screened negative for neurological and psychiatric disease. Five participants were excluded due to excessive motion. Twenty participants were included in the analysis in Experiment 1 (10 females) and 10 in Experiment 2 (5 females). Participants consented to the procedures in accordance with the Institutional Review Board of the University of California, Irvine, and received monetary compensation for their participation.

Stimuli

Experiment 1: Distinct stimulus set

In Experiment 1, the stimulus set consisted of two time-lapse videos (clips from Timestorm Films: https://vimeo.com/93003441) and two object pairs (Fig. 1A).

Experiment 2: Low-level image matching

In Experiment 2, the two contexts consisted of grayscale images (600×600 pixels) of Saint Peter's Basilica and the U.S. Capitol Building and the objects consisted of grayscale images (256×256 pixels) of car and house keys (Fig. 2A). We used a combined approach of image manipulation and model testing to diminish the presence of category information from the low-level visual features. First, we used the SHINE toolbox (Willenbockel et al., 2010) to equate luminance histograms across all of the scene stimuli and across all of the object stimuli. Second, we used a modeling approach to select images that were devoid of low-level category features.

For our scene images, similar to Marchette et al. (2015), we used pixel-wise correlation, the GIST computational model (Oliva and Torralba, 2001), and the HMAX computational model (two variants, one that used all images from the Fifteen Scene Categories dataset (Lazebnik et al., 2006) as prototypes and one that used a superset of our scene images as prototypes; we used the model from: Theriault et al., 2011). Additionally, similar to Kriegeskorte et al. (2008a), we used two models of V1 (one that included both simple and complex cells from HMAX and another that included only complex cells; Theriault et al., 2011), low-pass pixel-wise correlation (low frequency image features), high-pass pixel-wise correlation (high frequency image features), and Radon transform. We iteratively looped over a superset of our scene images and selected images for which all nine models showed no sign of a relationship between the scene images and the context matrix for both the selected stimulus set (40×40 matrix with 780 unique entries) and across the odd/even split (20×20 matrix with 400 unique entries; $-0.012 <$ Spearman's rank correlation < 0.011, all p's > 0.77). As a final control, we simulated an object being presented on top of each scene image by placing a black square (256×256 pixels) at the center of the image; importantly, similar results were obtained using this method.

For the object images, we used the same nine models as well as binary-silhouette correlation (similar to Kriegeskorte et al., 2008a).

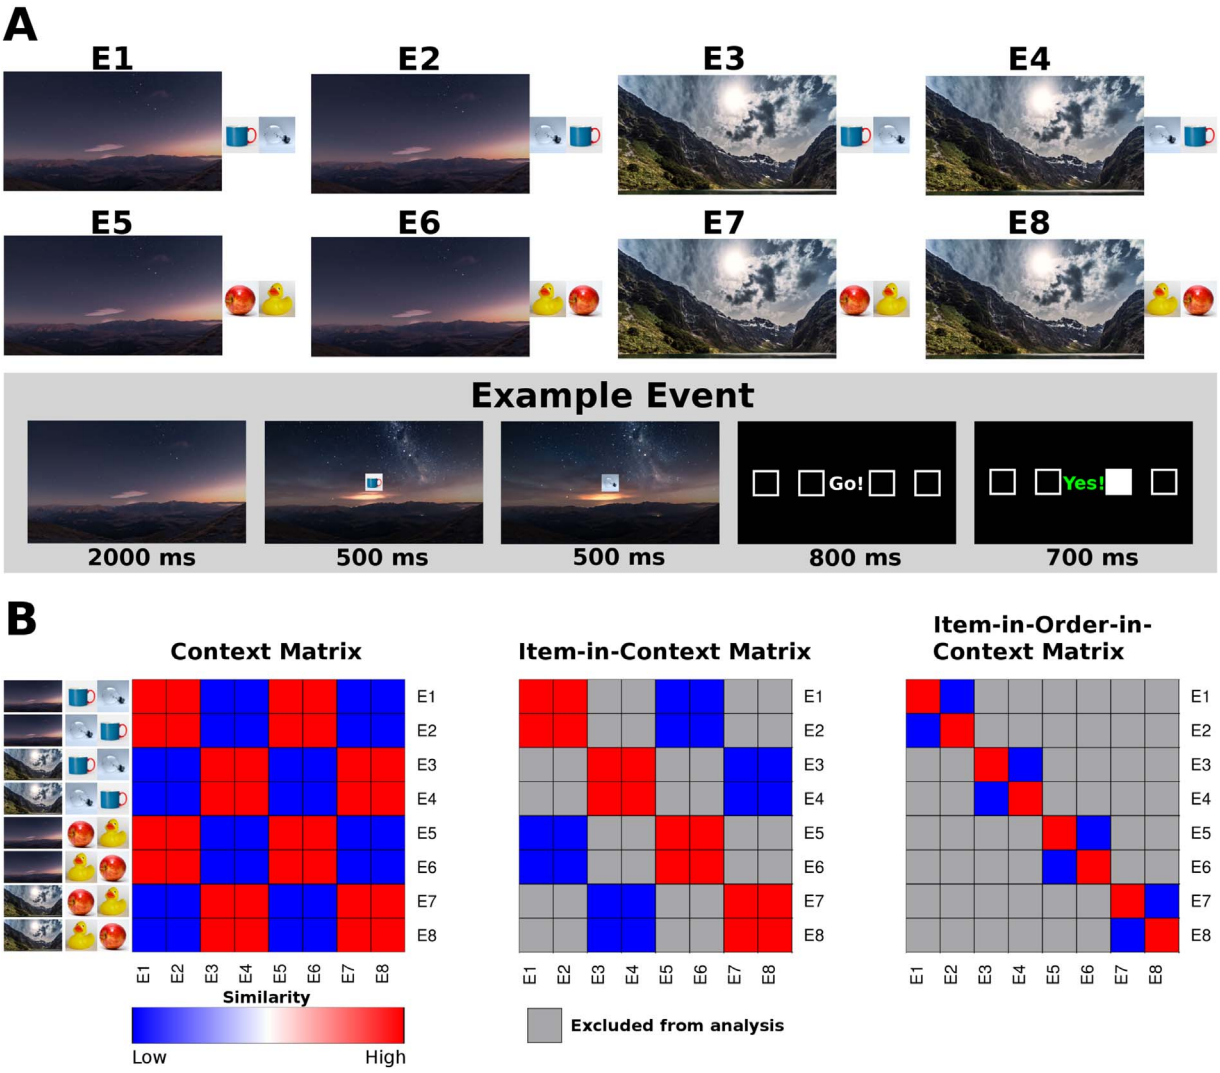

Fig. 1. Experiment 1 stimuli, event design, and model matrices. A) The task stimuli and an example event. The stimulus set consisted of two time-lapse videos (clips from Timestorm Films: https://vimeo.com/93003441) and two object pairs. Each event began with a 2000 ms presentation of a time-lapse video (depicted by the scene), then an object was displayed at the center of the video for 500 ms (depicted by the object on the left), which was then replaced by a second object which was displayed for 500 ms (depicted by the object on the right). Participants learned event-location associations the day prior to scanning, and were tested on the associations during scanning. B) Model matrices for our representational similarity analysis.

Similar to the analysis of the scene images, we used two versions of the HMAX model (one that used all images from the 256 Object Categories dataset (Griffin et al., 2007) as prototypes and one that used a superset of our object images as prototypes; Theriault et al., 2011). We iteratively looped over a superset of our object images and selected images for which all ten models showed no sign of a relationship to the object matrix for both the selected stimulus set (40×40 matrix with 780 unique entries) and across the odd/even split (20×20 matrix with 400 unique entries; $-0.016 \leq$ Spearman's rank correlation ≤ 0.014, all p's > 0.74).

Behavioral tasks

Experiment 1: Pre-scan training task

Participants were trained on an associative memory task one day prior to their scan session. The behavioral task was created using custom-written code and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Each event began with a 2000 ms presentation of a time-lapse video. An object was then displayed at the center of the video for 500 ms, which was replaced by a second object which was displayed for 500 ms, resulting in a 3 s event duration. One difference from the rodent task was that rather than using a physical reward, participants learned stimulus-location associations (the locations consisted of four squares on the computer screen, Fig. 1A) through trial and error learning with simple correct/incorrect feedback, similar to previous experiments in our laboratory (Law et al., 2005). In Experiment 1, the correct location was unique based on: 1) the video that was displayed, 2) the objects that were displayed, and 3) the order in which the objects were presented. Thus, the task required participants to make distinct responses for events that contain overlapping features. Participants responded using the 4 fingers on their right hand, and the event-location contingency was balanced across the two contexts (also see Supplementary Fig. 1). We used a response window of 800 ms, followed by 700 ms of feedback ("Yes!", "No!", or "?"=no response). The interstimulus interval consisted of a 400 ms fixation cross, a 700 ms arrow presentation (to which participants were instructed to indicate via button press whether it was pointing to the left or the right, which served as a non-mnemonic component of the interstimulus interval), and a 400 ms fixation cross, resulting in a trial length of 6 s (event duration=3 s, response and feedback=1.5 s, and interstimulus interval=1.5 s). We also included self-paced perceptual baseline trials (5.6 s blocks followed by a 400 ms fixation cross) in which participants were instructed to indicate as quickly and as accurately as possible which of four static noise boxes was the brightest (Law et al., 2005). The brightness of the target box was continuously titrated to maintain performance between 40 and 60% correct (chance=25%).

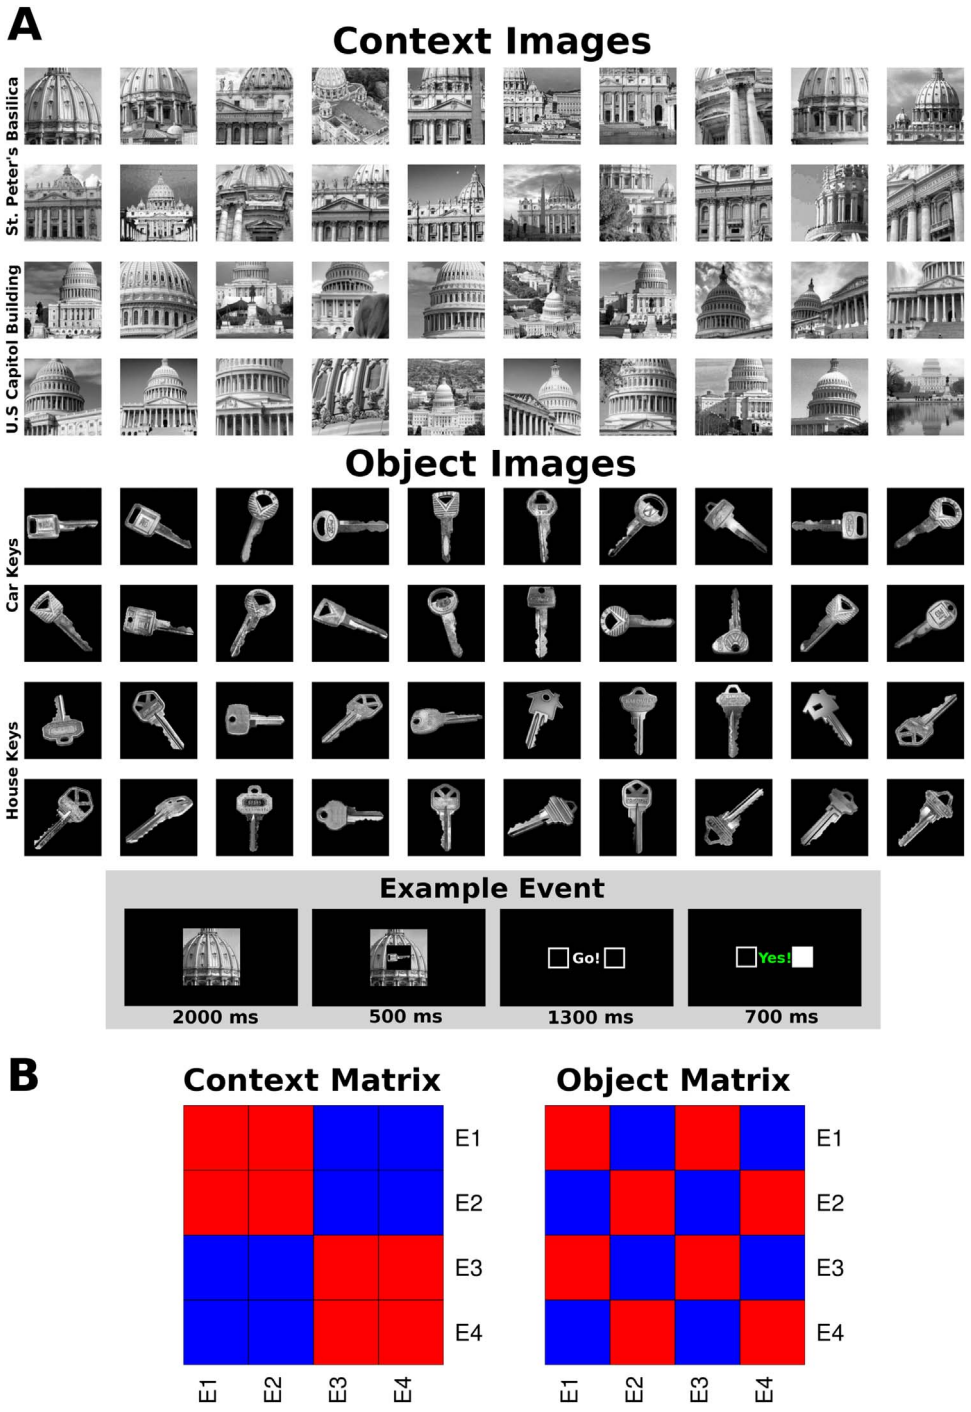

Fig. 2. Experiment 2 stimuli, event design, and model matrices. A) Task stimuli and an example event. The stimulus set consisted of grayscale images of Saint Peter's Basilica, the U.S. Capitol Building, car keys, and house keys. The odd numbered rows were used in odd runs of the task while the even numbered rows were used in even runs of the task. Each event began with a 2000 ms presentation of a context image, then an object was displayed at the center of the scene image of 500 ms. B) Model matrices for the representational similarity analysis. E1=St. Peter's Basilica+Car Key; E2=St. Peter's Basilica + House Key; E3=U.S. Capitol Building + Car Key; E4=U.S. Capitol Building + House Key.

At the beginning of the experiment, two events were presented. After a participant learned an event-location association to criterion—correct responses on 5 out of the last 6 responses—a new event was added to the unlearned queue, and the learned event was moved into the learned queue. Events in the learned queue were presented with p=0.3, while events in the unlearned queue were presented with p=0.7. Thus, participants continued to be tested on "learned" event-location pairs. We counterbalanced the order in which events were added to the unlearned queue. The session terminated after participants learned all 8 events to criterion.

Experiment 2: Pre-scan training tasks

Experiment 2 differed from Experiment 1 in the choice of stimuli and in the removal of the order component of the associative memory. To ensure that participants could readily discriminate between the "context" images, they were pre-trained on a category discrimination task for the 40 scene images. On each trial an image of either Saint Peter's Basilica or the U.S. Capitol Building was displayed for 1000 ms. Then, the image was removed and two boxes were displayed with text labels above each box ("St. Peter's Basilica" and "U.S. Capitol Building"; the left/right assignment of the text labels was random on each trial). We used a 1300 ms

response window, followed by 700 ms of feedback. The task terminated after participants learned each of the 40 scene images to criterion (correct responses on 5 out of the last 6 responses). In contrast to the associative memory task, only unlearned items were presented. Next, to ensure that participants could readily discriminate between the "object" images, participants performed an object discrimination task for the car and house key images. The task was identical to the scene task, except the text labels were "car" and "house."

After participants learned both the scene categories and the object categories to criterion, they then learned object-location associations. The event structure was similar to Experiment 1, with the exception that we used only one object per event (no item-order component), resulting in a stimulus duration of 2500 ms. Events were mapped to two responses (left versus right) and participants used their index finger to respond. Given the 500 ms reduction of stimulus presentation, we extended the response window from 800 ms to 1300 ms. We reduced the perceptual baseline task to contain two boxes, and performance was continuously titrated to maintain performance between 60 and 70% correct.

Experiment 1: fMRI task

Participants returned for their functional magnetic resonance imaging scan session one day after the training phase. During acquisition of structural scans, participants performed a warm-up phase, in which they were re-exposed to the associative memory task. The warm-up phase was included to attenuate novelty effects during the initial presentations of each event (Law et al., 2005). Once again, participants were initially tested on two event-location pairs. After participants learned an event-location association to criterion—two correct responses in a row—a new event was added to the unlearned queue and the learned event was added to the learned queue. As in the initial learning phase, items in the learned queue were presented with p=0.3, and the warm-up phase terminated after participants relearned all 8 events to criterion.

We designed the training and imaging paradigm to be similar to that used by McKenzie et al. (2014), in which they trained rats to criterion prior to neural recording. Thus, both studies investigated well- learned representations. During functional runs, participants were repeatedly tested on the event- location association task. Functional runs consisted of 4 presentations of each event as well as 5 self-paced perceptual baseline trials. We randomized the order of events within each run, with the exception that every run ended with one perceptual baseline trial to allow the hemodynamic response of the 2nd to last trial of the run to approach baseline prior to run completion. Participants completed 16 runs, resulting in 64 presentations of each event during functional scanning.

Experiment 2: fMRI task

There were a few minor differences between the fMRI task in Experiment 1 and Experiment 2. First, participants were given a reminder session for the stimulus categories (i.e., Saint Peter's Basilica, U.S. Capitol Building, car keys, house keys). Participants viewed a text label of the stimulus category, followed by a one second presentation of every image from the category (1 s presentation, 500 ms interstimulus interval). Participants saw each category two times. Second, functional runs consisted of 5 presentations of each event as well as 6 self-paced perceptual baseline trials. Third, de Bruijn sequences were used for stimulus ordering (Aguirre et al., 2011). A unique sequence was used for each run (randomized across subjects), and we selected sequences that ended with perceptual baseline trials to allow the hemodynamic response of the 2nd to last trial of the run to approach baseline prior to run completion. Carry-over sequences, such as de Bruijn sequences, match the number of times that each stimulus precedes every other stimulus, thus controlling for stimulus carry-over effects and theoretically increasing the detection power in between-run pattern analysis (Aguirre, 2007; Aguirre et al., 2011). The length of carry-over sequences was prohibitively large for Experiment 1, but the reduction of stimulus conditions allowed an entire de Bruijn sequence to be presented within a short run (for an argument for using short runs for pattern analysis see: Coutanche and Thompson-Schill, 2012; Davis and Poldrack, 2013). Finally, participants completed 12 runs, resulting in 60 presentations of each event during functional scanning.

MRI data acquisition

Data were acquired from a 3.0-T Philips scanner, using a 32 channel sensitivity encoding (SENSE) coil at the Neuroscience Imaging Center at University of California, Irvine. A high-resolution 3D magnetization-prepared rapid gradient echo (MP-RAGE) structural scan (0.75 mm isotropic) was acquired for each participant. Functional MRI scans consisted of a T2*-weighted echo planar imaging sequence using blood-oxygenation-level-dependent contrast (BOLD; repetition time [TR]=2500 ms, echo time=26 ms, flip angle=70 degrees, 46 slices, 2.5×2.5 mm in plane resolution, 2.3 mm slice thickness with a 0.2 mm gap). Each functional run was padded with an initial 4 "dummy" dynamics, which were immediately discarded to ensure T1 stabilization. In Experiment 1, 90 dynamics were collected per run and 16 functional runs were collected for each participant; however, the 16th run for one participant was not analyzed due to large between-run motion. In Experiment 2, 64 dynamics were collected per run and 12 functional runs were collected for each participant.

fMRI data preprocessing

Data were preprocessed using Analysis of Functional NeuroImages (AFNI; Cox, 1996). Functional MRI data were motion corrected using rigid-body transformation using the function align_epi_anat.py (Saad et al., 2009). Data were quadratically detrended and high pass filtered ($f > 0.01$ Hz), using the 3dBandpass function. To preserve fine-grained information, the data were left unsmoothed. We manually defined the hippocampus, PHC, and PRC on a custom template brain according to previously defined landmarks (for more details see: Law et al., 2005). We used Advanced Normalization Tools (ANTs; Avants et al., 2008) to warp each individual participant's MP-RAGE structural scan into our custom template space. The inverse warp vectors were used to create masks for the hippocampus, PHC, and PRC within each participant's original space (Huffman and Stark, 2014). Freesurfer's isthmus cingulate label (Desikan et al., 2006) was used to define RSC/PCC. The isthmus cingulate mask contains voxels from RSC (traditionally defined as Brodmann's areas 29 and 30; Vann et al., 2009) and a portion of PCC caudal to RSC. Control regions, left primary motor cortex and bilateral V1, were generated using Freesurfer's precentral gyrus label (Destrieux et al., 2010) and Freesufer's V1 atlas (Hinds et al., 2008), respectively. Masks were resampled to 2.5 mm isotropic (the fMRI grid) and further masked to contain completely-sampled voxels. We used a combined anatomical and functional approach to define parahippocampal place area (PPA) and retrosplenial complex (RS-Complex; Julian et al., 2012). We warped the anatomical masks to each subject's native space and selected the 100 most active voxels (all events versus perceptual baseline) in each anatomical mask in each hemisphere and merged the resultant files to create bilateral masks for each ROI (Marchette et al., 2015; Vass and Epstein, 2016).

Representational similarity analysis

Data were analyzed using AFNI, custom-written code in Python and R, and PyMVPA (Hanke et al., 2009) on a GNU/Linux platform using the NeuroDebian package repository (Hanke and Halchenko, 2012). The fMRI data were split in half—odd and even runs—and a block-based general linear model (GLM; AFNI's 3dREMLfit function; the block length was 3 s in Experiment 1 and 2.5 s in Experiment 2—i.e., the length of the event presentations) was used to generate beta values in each voxel. During volume registration, 6 motion parameters were generated (3 translation parameters and 3 rotation parameters).

Framewise displacement is defined as the sum of the absolute value of the difference between each of the 6 motion parameters (3 translation, 3 rotation: we approximated that 1 degree of rotation caused 1 mm of movement) between successive frames (Power et al., 2012). Frames with framewise displacement exceeding 0.5 mm and 1 frame before and 2 frames after were censored from the analysis (for a similar approach to functional connectivity analysis see: Power et al., 2012). Within each split, the mean pattern of activity across all events was subtracted from each event-specific beta vector (Haxby et al., 2001). Pearson's correlation coefficients were calculated between each event-specific beta vector across the split-halves, resulting in non-symmetrical representational similarity matrices. Hypothesis-driven analyses were conducted by calculating Spearman's rank correlation coefficient between each participant's representational similarity matrix and predefined matrices (Figs. 1B and 2B).

In Experiment 1, we performed an iterative approach in our correlation analysis between each subject's representational similarity matrix and our model matrices. We chose this approach for two reasons: 1) the model matrices were correlated with each other, thus precluding analysis within a single model, 2) correlation analysis is not sensitive to the magnitude of the values within the similarity matrix (as opposed to comparing whether within-category correlations were numerically larger than between-category correlations), thus allowing us to examine the pattern of similarity regardless of the magnitude of the values (Kriegeskorte et al., 2008a). In Experiment 1, we began with the model matrix on the left side of Fig. 1B (i.e., the context matrix) and proceeded rightwards only for matrices that were significantly related to the model matrix in the present step (i.e., we terminated analysis for an ROI when the relationship between the ROI matrix and model matrix failed to reach significance). A previous study Kriegeskorte et al. (2008b) investigated the cross-species correlation between portions of representational similarity matrices, which is similar to our approach of comparing portions of representational similarity matrices to model matrices. In both experiments, we used Spearman's rank correlation, rather than Pearson's correlation coefficient, because it is better suited for investigating the relationship between ordinal models and representational similarity matrices (Kriegeskorte et al., 2008a).

Permutation analysis

Previous reports have suggested that nonparametric methods are preferable to classical statistical tests for analyzing the significance of the relationship between representational similarity matrices (Kriegeskorte et al., 2008a, 2008b); therefore, we used a two-step permutation method to determine statistical significance (for related approaches to classification analysis see: Chen et al., 2011; Liang et al., 2013; Stelzer et al., 2013; Etzel, 2015). For each participant, the empirical similarity matrix was randomly shuffled and we calculated the Spearman's rank correlation between the resultant matrix and the intact model matrix. The resultant value was then Fisher's r-to-z transformed (z[r]), using the inverse hyperbolic tangent function. We performed this process 10,000 times for each participant to generate null distributions at the subject level. To maintain similarity to the empirical analysis, we used the same permutation of the labels across participants (Etzel, 2015). The null distributions were averaged across participants to generate a group-level null distribution. Two-tailed nonparametric p-values were calculated using the following equation (Ernst, 2004):

$$p = \frac{1 + \sum_{i=1}^{10,000} I\left(|t_i - \bar{t}| \geq |t^* - \bar{t}|\right)}{1 + 10,000} \quad (1)$$

where $I(\cdot)$ is the indicator function which sets the value to 1 if the statement is true and to 0 otherwise, t is our test statistic (z[r]), t_i represents the ith value of the permutation vector, $\bar{t}$ represents the mean of the permutation vector (mean of the null distribution), and t^* represents the empirical (observed) mean. This calculates the probability that a null value was at least as far (in both directions) from the mean of the null distribution as the empirical value (akin to a two-tailed test). To maximize the degree of similarity to a full permutation approach, 1 is added to both the numerator and the denominator of the equation (i.e., in a full permutation the stimulus labels would be in the correct order exactly once, hence the lowest p-value attainable is 1 divided by the number of combinations). For Experiment 2, the full permutations were tractable, thus p values were calculated using the following equation (Ernst, 2004):

$$p = \frac{\sum_{i=1}^{N} I\left(|t_i - \bar{t}| \geq |t^* - \bar{t}|\right)}{N} \quad (2)$$

where N is the total number of combinations. In Experiment 2, we performed nonparametric difference tests by subtracting permutation matrices from each other across ROIs. The resultant permutation difference matrices were averaged across participants and significance was assessed using Eq. (2), where the t's represent difference values. We compared the results of Experiment 1 and 2 using a permutation analysis in which we shuffled the experiment labels and calculated the mean difference between the shuffled groups. We used Eq. (1) to calculate p-values based on the empirical difference in means relative to 10,000 random permutations.

Informational correlativity analysis

In Experiment 1, we tested the hypothesis that the hippocampus, PHC, and RSC/PCC contain similar information about context on a trial-by-trial basis, using a variant of "informational connectivity" (Coutanche and Thompson-Schill, 2013; Huffman and Stark, 2014), which we refer to here as informational correlativity. We used an extension of the LS2 procedure (Turner et al., 2012) to obtain individual trial estimates of activity. Briefly, for each trial, we ran a block-based GLM analysis (using AFNI's 3dDeconvolve function) that included an individual trial regressor and 8 event-specific regressors that coded for every other trial. We then performed 512 iterations of this procedure (i.e., 8 events×64 presentations of each event). Each iteration of the GLM incorporated the same censor vectors as before (i.e., framewise displacement >0.5 mm); additionally, to mitigate the adverse effect of noisy individual trial estimates from subsequent analysis, we removed trials that had a motion event within approximately 15 s of the onset of the trial (the exact duration was variable because the image acquisition was not time-locked to the trial presentation).

To maintain similarity to the initial analysis, we used a split-halves approach in this procedure. Similarly, we subtracted the mean pattern of activity across all events from each event-specific beta vector, separately within each split. We then averaged the patterns of activity within each context (i.e., E1, E2, E5, and E6 were averaged to create an average "context 1" pattern of activity and E3, E4, E7, and E8 were averaged to create an average "context 2" pattern of activity) separately within each split. For each trial, we calculated z[r] Pearson's correlation coefficient between the pattern of activity on that trial and the average context 1 and context 2 patterns of activity from the other split. The value for each trial was set to the correlation to the same context minus the correlation to the other context (Coutanche and Thompson-Schill, 2013). Thus, values greater than 0 denote correct "neural discrimination" between the two contexts and the distance from 0 provides an index of discriminability on that trial. We refer to the trial-by-trial vector of such values to as the multivariate pattern discriminability trial-series (for a related approach to functional connectivity see: Rissman et al., 2004). For each participant, we calculated z[r] Spearman's rank correlation between the multivariate pattern discriminability trial-series in the hippocampus, PHC, and RSC/PCC. We

averaged the resultant values to obtain the empirical group mean z[r] Spearman's rank correlations. To assess significance, we used a permutation approach in which we randomized the order of one of the ROI's multivariate pattern discriminability trial-series within each run. We used within-run permutations, rather than permuting the entire trial-series, to mitigate the possibility that between-run differences would artificially reduce the permuted correlations. We then calculated z[r] Spearman's rank correlation between the randomized ROI's vector and the intact ROI's vector. This procedure was carried out 10,000 times per subject. Group-level permutation analysis was conducted by averaging the permuted distributions across subjects, and p values were obtained using Eq. (1). In this application, the permutations were independent across subjects because the trial order was independent across subjects.

Multidimensional scaling analysis

We performed multidimensional scaling (MDS), which is a data-driven, data-reduction approach that allows visualization of the major components of similarity matrices such as those obtained using representational similarity analysis on fMRI data (Kriegeskorte et al., 2008a, 2008b). We first generated a symmetrical representational similarity matrix by averaging values across the diagonal of the matrix (i.e., the same pairs of events across splits). We then converted the matrix from z[r] Pearson's correlation coefficient to Pearson's correlation coefficient (r), and converted the resultant matrix to correlation distance (1−r). We extracted the lower triangle of the correlation distance matrix and performed multidimensional scaling using the criterion of metric stress (Kriegeskorte et al., 2008a, 2008b) using package cmdscale in R. We used custom-written code to place the event stimuli at the coordinates calculated by the multidimensional scaling procedure.

Relationship between representations and behavioral performance on the associative memory task

In Experiment 1, we investigated the relationship between representations and behavioral performance. For each participant, we defined model fit as the z[r] Spearman's rank correlation between their similarity matrix and our proposed model. We calculated the proportion of correct responses during functional scanning within each participant, excluding trials in which the participant did not respond within the response window. We calculated Spearman's rank correlation coefficient between model fit and proportion correct. We used Spearman's rank correlation because it does not require the assumption that the two variables are normally distributed (as opposed to Pearson's correlation). To assess significance, we calculated the t statistic using the following equation (Krzanowski, 2000):

$$t = r\sqrt{\frac{k-1}{1-r^2}} \tag{3}$$

where r is Spearman's rank correlation, and k is equal to n−1. We obtained a p value from Student's t distribution with k−1 degrees of freedom (20-1-1=18).

To mitigate the possibility of a spurious effect of head motion (Power et al., 2012) on the observed relationship between behavioral performance and model fit, we performed a follow-up analysis using a partial correlation approach. Specifically, we examined the relationship between model fit and behavioral performance while holding the effect of head motion constant. Previous reports have used mean motion—defined as the mean amount of motion between successive frames based on the sum of the 3 translation parameters—as a measure of head motion (Van Dijk et al., 2012). Mean motion has previously been shown to be strongly correlated with the total number of motion events and it has been shown to be a reliable index of subject-specific head motion (Van Dijk et al., 2012). We used a related measure, mean framewise displacement, as our measure of head motion. As mentioned above, framewise displacement is the sum of motion across all 6 alignment parameters (3 translation and 3 rotation parameters) between successive frames (Power et al., 2012). To calculate the partial Spearman's rank correlation between the model fit and the proportion correct while holding the effect of head motion constant, we used the following equation (Krzanowski, 2000):

$$r_{xy.z} = \frac{r_{xy} - r_{xz}r_{yz}}{\sqrt{\left(1-r_{xz}^2\right)\left(1-r_{yz}^2\right)}} \tag{4}$$

where x represents the model fit array, y represents the proportion correct array, and z represents the mean framewise displacement array. In our application, r represents Spearman's rank correlation. The t statistic was calculated using Eq. (3). With one variable held constant (i.e., z), k=n−1−1. Therefore, the statistical test for partial correlation is the same as in typical correlation analysis with fewer degrees of freedom (Krzanowski, 2000).

Whole-brain searchlight analysis

We performed a whole-brain searchlight analysis using a searchlight radius of 3 voxels within the *sphere_searchlight* function in PyMVPA. We ran the analysis in native space and warped the results to our group template using ANTs (Avants et al., 2008). For each contrast, we used a voxel-wise threshold of $p < 0.01$ (parametric) and a cluster threshold of $p < 0.05$ (cluster threshold was determined using Monte Carlo simulation with a simulated blur of 3.75 mm FWHM, i.e., half of the searchlight radius). First, we investigated the relationship to the context matrix. Next, we masked the results for the item-in-context matrix by the regions that were significantly related to the context matrix. Finally, we masked the results for the item-in-order-in-context matrix by the regions that were significantly related to both the context matrix and the item-in-context matrix. The overlap map was warped to an inflated brain for visualization using FreeSurfer.

Relationship between the empirical V1 similarity matrix and the model V1 similarity matrix

In Experiment 2, we investigated the similarity of the empirical V1 similarity matrix and the model V1 similarity matrix (using HMAX; Theriault et al., 2011). We conducted a trialwise analysis in which we modeled activity in response to both the scene image and the object image on each trial. Specifically, the V1 model was “shown” the same stimulus sequence as the participant. On each trial, the V1 model response vectors from Layer 2 of HMAX were extracted in response to the scene image and in response to the scene and object images. The resultant vectors were combined using weighted averaging (4/5 of the scene only vector plus 1/5 of the scene plus object vector—i.e., 2000 ms scene presentation and 500 ms scene plus object presentation). The mean pattern of activity was removed within the odd and even splits, and correlation matrices were generated by correlating the pattern of activity across all of trials across the odd and even splits. A similar approach was conducted to generate the empirical similarity matrix using the single trial estimation techniques described for representational correlativity analysis. We calculated Spearman's rank correlation between the trialwise empirical V1 similarity matrix and the model V1 similarity matrix using the motion censoring steps described for representational correlativity analysis.

Bayes factor analysis

In Experiment 2, we used a Bayes factor analysis to examine whether the data were more consistent with the null hypothesis (mean z[r] Spearman's rank correlation=0) or the alternative hypothesis (mean z[r] Spearman's rank correlation > 0), where BF_{01} denotes the

odds ratio of the null hypothesis relative to the alternative hypothesis (Rouder et al., 2009). This analysis is a one-sample test in which the observations are comprised of each participant's z[r] Spearman's rank correlation between the similarity matrix in a given ROI and a hypothesis matrix (e.g., the context matrix). As recommended by Rouder et al. (2009), we used JZS Bayes factors with scaling factor r=1, within the *BayesFactor* package in R.

Results

Experiment 1

The task in Experiment 1 consisted of eight events, which differed in terms of the time-lapse video that was displayed ("context"), the items that were displayed ("item"), and the order in which the items were presented ("order"; Fig. 1A). Participants (n=20) learned event-location associations the day prior to fMRI scanning. The correct location depended on the video that was displayed, the objects that were displayed, and the order in which the objects were presented. Participants were trained to criterion the day prior to their scan session. During scanning, participants were repeatedly tested on each event-location association, during which they continued to exhibit above-chance performance (mean proportion correct=0.86, t_{19}=27.15, $p < 0.0001$, 95% CI [0.81, 0.90]). Similar to McKenzie et al. (2014) and Keene et al. (2016), we performed representational similarity analysis (Kriegeskorte et al., 2008a, 2008b), in which we calculated Pearson's correlation coefficient between patterns of activity in response to each of the 8 events across the odd and even runs.

Investigation of the representation of context

To test the hypothesis that the hippocampus, PHC, and RSC/PCC are involved in context representation, we generated a context matrix that contains uniformly large values for events that share the same context and uniformly small values for events that contain different contexts (Fig. 1B). We calculated Spearman's rank correlation coefficient (Fisher's r-to-z transformed: z[r]) between each participant's similarity matrix and the context matrix. We used group-level two-tailed nonparametric p values (Ernst, 2004) to assess significance (for related one-tailed approaches to classification analysis see: Chen et al., 2011; Liang et al., 2013; Stelzer et al., 2013; Etzel, 2015). The PHC and RSC/PCC similarity matrices were significantly related to the context matrix (PHC: mean z[r] Spearman's rank correlation [M]=0.15; RSC/PCC: M=0.81; both $p < 0.0001$; Fig. 3A and B) but a relationship was not observed for the hippocampus similarity matrix (M=0.012, p=0.71). We also tested the hypothesis that PRC carries object information; however, the PRC similarity matrix failed to exhibit a significant relationship to either the object matrix (M=0.032, p=0.30) or the context matrix (M=0.032, p=0.32).

We performed a searchlight analysis within the hippocampus to investigate whether there was a relationship to the context matrix in a portion of the hippocampus, which may have been obscured by uninformative voxels. Using a searchlight radius of 3 voxels, we observed a significant cluster in the left posterior hippocampus (33 voxels, parametric voxel-wise threshold $p < 0.05$; Fig. 3C). A follow-up analysis, in which we warped the searchlight cluster mask to each participant's native space and ran a region of interest (ROI) analysis, indicated that the cluster itself was significantly related to the context matrix (p=0.017). The follow-up analysis is circular but it is required to conclude that the cluster itself is informative (Etzel et al., 2013); therefore, these results bolster the conclusion that the hippocampal cluster is related to the context matrix.

To eliminate the possibility that the context results were driven solely by strong relationships between an event and "itself"—i.e., because the identity matrix is weakly correlated to the context matrix—we performed a control analysis that excluded the entries from the main diagonal of the matrix. The results remained in PHC (M=0.15; $p < 0.0001$), RSC/PCC (M=0.78; $p < 0.0001$) and in the hippocampal searchlight cluster ($p < 0.01$). Conversely, as a control region, the left precentral gyrus (primary motor cortex) similarity matrix failed to exhibit any sign of a relationship (M=0.0072; p=0.85); instead (and as one might expect), it was related to the correct-response matrix (Supplementary Fig. 1). In contrast, there was no sign of a relationship between the PHC and RSC/PCC similarity matrices and the motoric-based correct-response matrix (Supplementary Fig. 1). These results eliminate the possibility that: 1) the relationship between the similarity matrices and the context matrix was driven by correlations between patterns of activity in response to an event and "itself", 2) the relationships to the context matrix were confounded by key-press differences between events in opposing contexts.

We used informational correlativity (Coutanche and Thompson-Schill, 2013; Huffman and Stark, 2014) to test the hypothesis that the hippocampal searchlight cluster, PHC, and RSC/PCC contain similar representations of context on a trial-by-trial basis and thus are related in their processing of contextual information. There was a significant relationship between trial-by-trial context representation in PHC and RSC/PCC (M=0.34, $p < 0.0001$) and both cortical regions and the hippocampal searchlight cluster (both M=0.085, $p < 0.0001$; Fig. 3D).

Investigation of the representation of items in context

To test the hypothesis that PHC and RSC/PCC contain item-in-context information, we calculated z[r] Spearman's rank correlation between each participant's similarity matrix and the item-in-context matrix. This analysis investigates whether events that share the same context and the same items are represented more similarly than events that share the same context but contain different items (all other events are excluded from the analysis; see Fig. 1B). The RSC/PCC similarity matrix was significantly related to the item-in-context matrix (M=0.14, $p < 0.01$; Fig. 3E) but this relationship was not observed for the PHC similarity matrix (M=0.012, p=0.76). In contrast, the relationship between the RSC/PCC similarity matrix and an item-out-of-context matrix (uniformly large values for events that shared the same items but different contexts and uniformly small values for events that contain different items and different contexts) failed to reach significance (M=0.047, p=0.32). To get a parsimonious account of the information coded for in the matrices and to visualize this, we performed MDS. The first two dimensions captured by MDS were context and items (Fig. 3F). Altogether, these results are consistent with the notion that RSC/PCC contains conjunctive item-in-context representations.

To extend this, we tested the hypothesis that representations in RSC/PCC are related to memory performance by generating a model of representations in RSC/PCC that contains context and item-in-context information (Fig. 3G). There was a significant relationship between model fit (z[r] Spearman's rank correlation between the RSC/PCC similarity matrix and our proposed model) and proportion correct (Spearman's rank correlation=0.51, t_{18}=2.51, $p < 0.05$; Fig. 3H), which maintained when controlling for the potential confound of head motion (partial Spearman's rank correlation=0.47, t_{17}=2.18, $p < 0.05$).

Investigation of parahippocampal place area and retrosplenial complex

There has been extensive evidence of scene and context processing in the parahippocampal place area (PPA) and the retrosplenial complex (RS-Complex; e.g., Epstein and Kanwisher, 1998; Epstein et al., 2007; Julian et al., 2012; Vass and Epstein, 2013, 2016; Marchette et al., 2014, 2015). These regions are in close anatomical proximity to PHC and RSC/PCC, however they are largely non-overlapping. Specifically, PPA tends to be located posterior to PHC (i.e., along the parahippocampal gyrus, but posterior to the landmarks like the splenium of the corpus callosum) and retrosplenial complex tends to be located posterior to RSC/PCC (Supplementary Fig. 2A). Within the anatomical masks for PPA and RS-Complex, we selected the 100 most active voxels (events > perceptual baseline; Marchette et al., 2015) from each hemisphere. We combined the

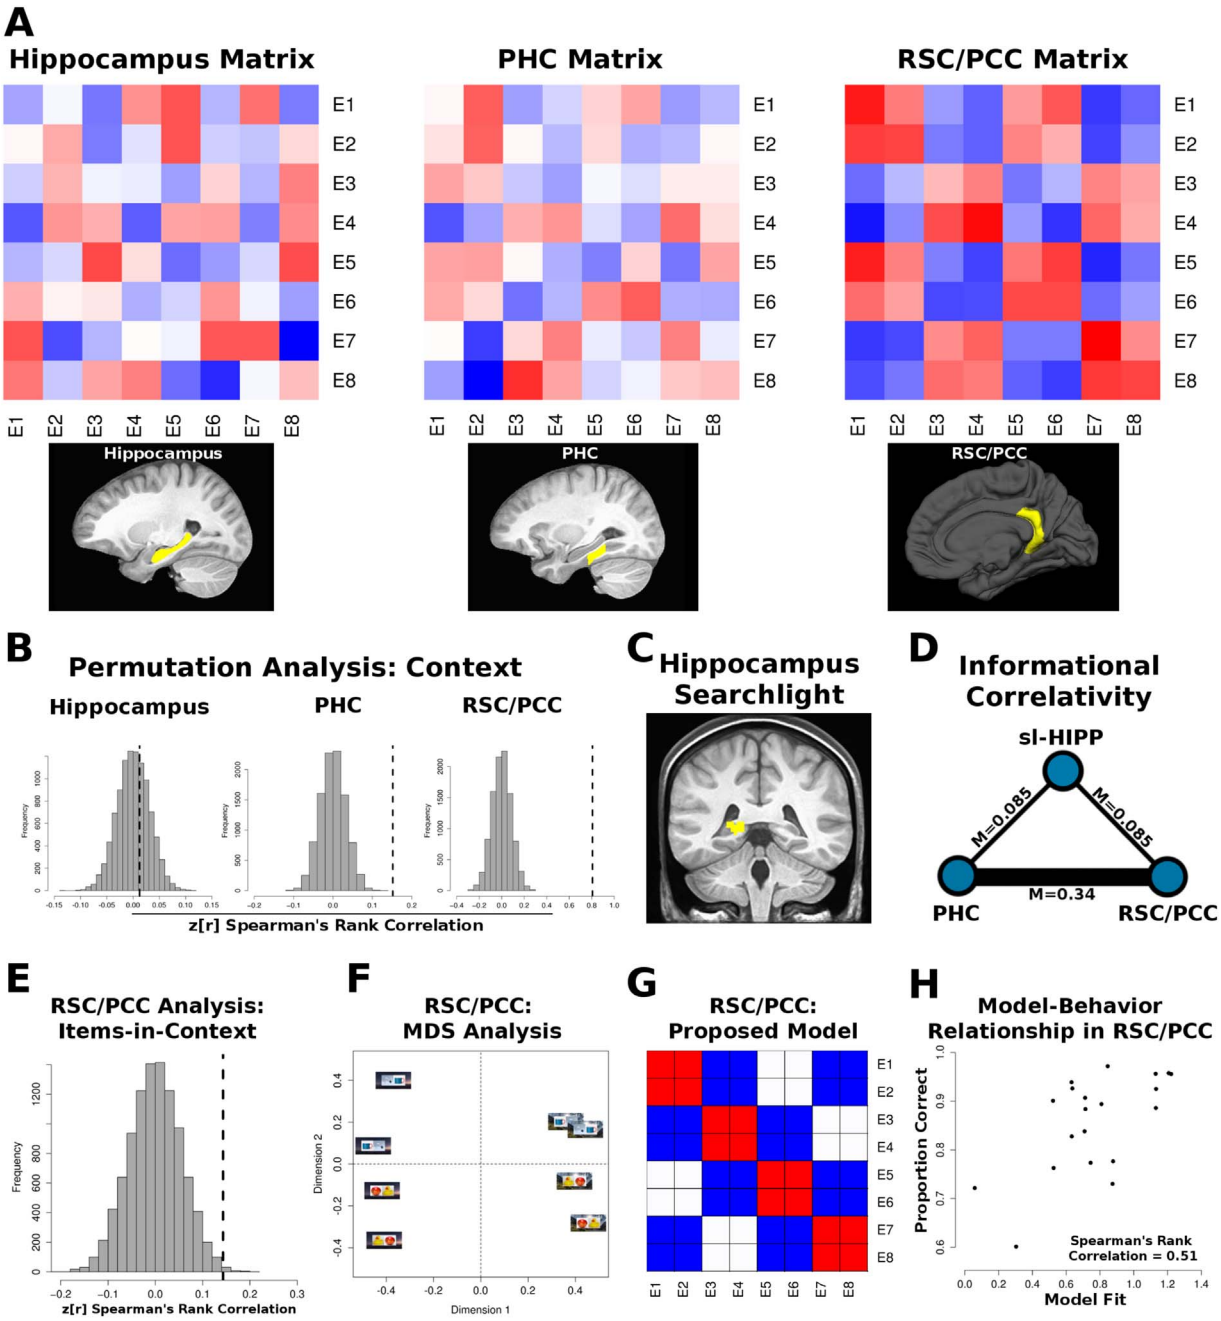

Fig. 3. Investigation of the hippocampus, PHC, and RSC/PCC. A) Average correlation matrices. B) Permutation analysis revealed a significant relationship between the context matrix and the PHC and RSC/PCC similarity matrices (p's < 0.0001) but the relationship failed to reach significance for the hippocampus similarity matrix (p=0.71). C) A searchlight analysis within the hippocampus revealed a significant cluster in the left posterior hippocampus. D) Informational correlativity analysis revealed a significant relationship between trial-by-trial representations in all three regions (all p's < 0.0001; sl-HIPP=hippocampus searchlight analysis cluster). E) Permutation analysis revealed a significant relationship between the RSC/PCC similarity matrix and the item-in-context matrix. F) Multidimensional scaling (MDS) analysis provided further evidence for item-in-context representations in RSC/PCC. The first dimension split based on event videos and the second dimension split based on event objects, which is consistent with the notion that RSC/PCC contains item-in-context information. G) Proposed model of RSC/PCC representations, which contains context and item-in-context information. H) There was a relationship between model fit in RSC/PCC and performance on the task (Spearman's rank correlation=0.51, t_{18}=2.51, p < 0.05).

masks into bilateral PPA and RS-Complex. The PPA and RS-Complex similarity matrices were significantly related to the context matrix (PPA: M=0.61, p < 0.0001; RS-Complex: M=1.04, p < 0.0001; Supplementary Fig. 2B) and the item-in-context matrix (PPA: M=0.15, p < 0.005; RS-Complex: M=0.13, p < 0.01; Supplementary Fig. 2C), similar to the results in RSC/PCC.

Investigation of V1

We investigated whether there were similar effects in primary visual cortex (V1), which would indicate the presence of a low-level visual confound. We found a similar pattern of results to RSC/PCC, including a strong relationship to both the context matrix (M=1.23, p < 0.0001) and the item-in-context matrix (M=0.16, p < 0.01). Similarly, the relationship to the item-out-of- context matrix failed to reach significance (M=0.014, p=0.81). We observed a relationship between the V1 similarity matrix and the item-in-order-in-context model (M=0.21, p < 0.005; Fig. 4A-B), which reveals that patterns of activity in response to events that share the same context, items, and order of item presentation are represented more similarly than events that share the same context and items but a switched order of item presentation. Similar to RSC/PCC, there was a

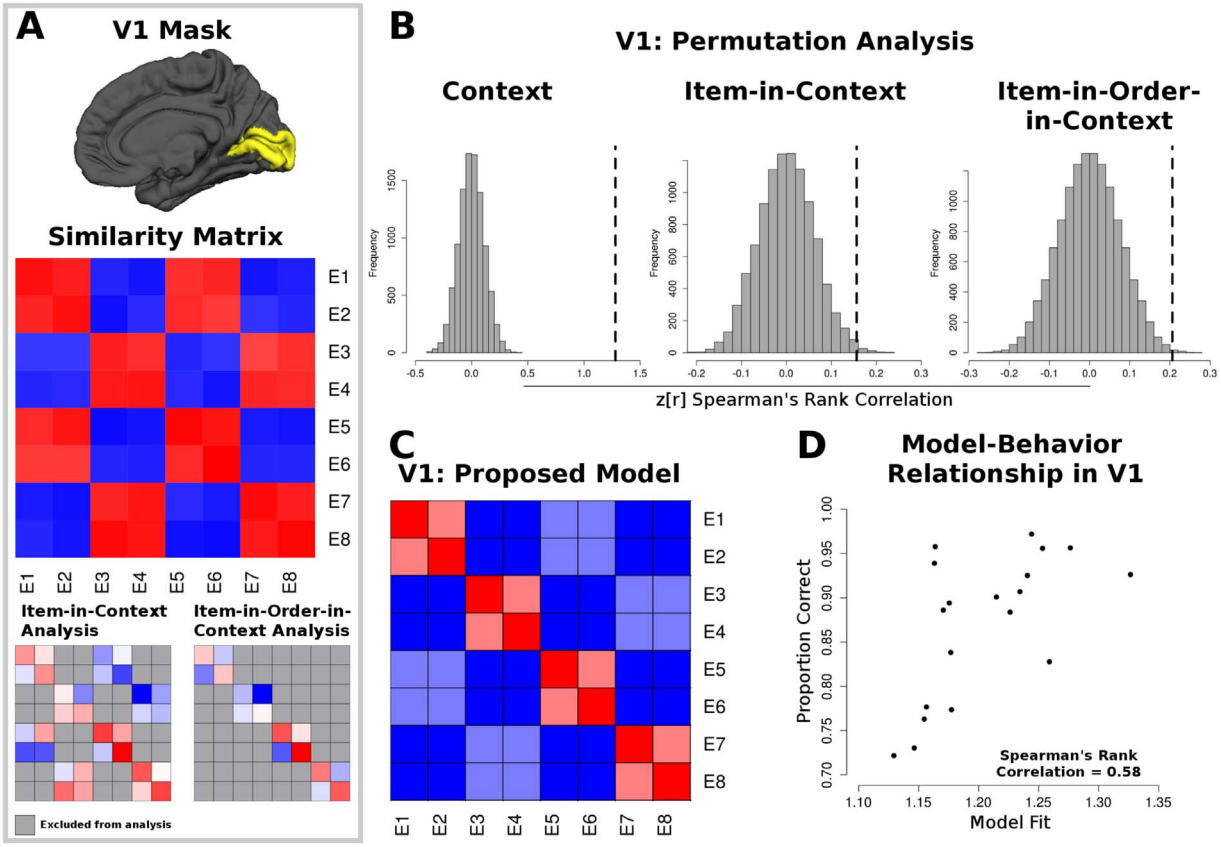

Fig. 4. The investigation in V1 revealed that our effects in Experiment 1 are confounded by the physical differences between stimuli. A) V1 mask (top), average correlation matrix (middle), the average correlation matrix of the data used for the item-in-context analysis (bottom left), and the average correlation matrix of the data used for the item-in-order-in-context analysis (bottom right). Note, the data were replotted in the bottom panel to reveal the visual similarity to both the item-in-context matrix and the item-in-order-in-context matrix (see Fig. 1B). B) Permutation analysis revealed a significant relationship between the V1 similarity matrix and the context matrix (M=1.23, $p < 0.0001$), the item-in-context matrix (M=0.16, $p < 0.01$), and the item-in-order-in-context matrix (M=0.21, $p < 0.005$). C) Proposed model of V1 representations, which contains context, item-in-context, and item-in-order-in-context information. D) There was a relationship between model fit and performance on the task (Spearman's rank correlation=0.58, $t_{17}=2.97$, $p < 0.01$).

significant relationship between model fit in V1 (the proposed model contained context, item-in-context, and item-in-order-in-context information; Fig. 4C) and proportion correct (Spearman's rank correlation=0.58, $t_{17}=2.97$, $p < 0.01$; Fig. 4D), which maintained when controlling for the effect of head motion (partial Spearman's rank correlation=0.58, $t_{16}=2.88$, p=0.01; one participant was dropped from the analysis due to low model fit, however, the effect was stronger, for both approaches, when the participant was included in the analysis).

Whole-brain searchlight analysis

To assess the extent of these findings, we used a whole-brain searchlight analysis (3 voxel radius) to investigate the prevalence of the relationship to each of our model matrices. The analysis revealed clusters in all four lobes, indicating this was not highly localized. However, the most prominent findings were in the occipital lobes, where a large proportion of occipital voxels were significantly related to the context matrix, the item-in-context matrix, and the item-in-order-in-context matrix (Supplementary Fig. 3). Altogether, the findings in V1 and the whole-brain searchlight analysis highlight the possibility that the context effects that we observed in the posterior hippocampus, PHC, and RSC/PCC were influenced by low-level visual differences between the stimuli that comprised the events. Therefore, our primary aim in Experiment 2 was to investigate whether our ROIs would still carry contextual information after eliminating low-level differences between our contexts.

Experiment 2

There are many possible approaches to reduce the influence of low-level features (see Discussion), but we chose to investigate whether our ROIs exhibit stable representations across multiple images of two contexts and objects. We used stimulus filtering and computational modeling to diminish the presence of low-level sensory differences between our contexts and objects. Participants (n=10) learned to discriminate between the two contexts (images of Saint Peter's Basilica and the U.S. Capitol Building) and the two objects (images of car keys and house keys), and then learned event-location associations (similar to Experiment 1, but without the order manipulation; Fig. 2A). Participants were trained to criterion on the discrimination tasks and the associative memory task to ensure that they could rapidly discriminate between our contexts and objects and to ensure that they could accurately perform the associative memory task. During scanning, participants were repeatedly tested on each event-location association, during which they continued to exhibit above-chance performance (mean proportion correct=0.86, $t_9=14.95$, $p < 0.0001$, 95% CI [0.80, 0.91]). Given that the nature of the memory task was unchanged, we hypothesized that regions that carry information about contexts and objects—as those terms relate to performance on this task—should do so in an invariant manner. Conversely, if our ROI-based results from Experiment 1 were influenced by low-level stimulus features, then we should fail to observe a relationship to the context matrix in our ROIs.

Investigation of in V1

First, we investigated whether our stimulus filtering and computational modeling approaches removed the low-level visual confound that was observed in Experiment 1. Importantly, the V1 similarity matrix showed no sign of a relationship to either the context matrix (M=−0.029, p=0.81) or the object matrix (M=0.0025, p=0.99;

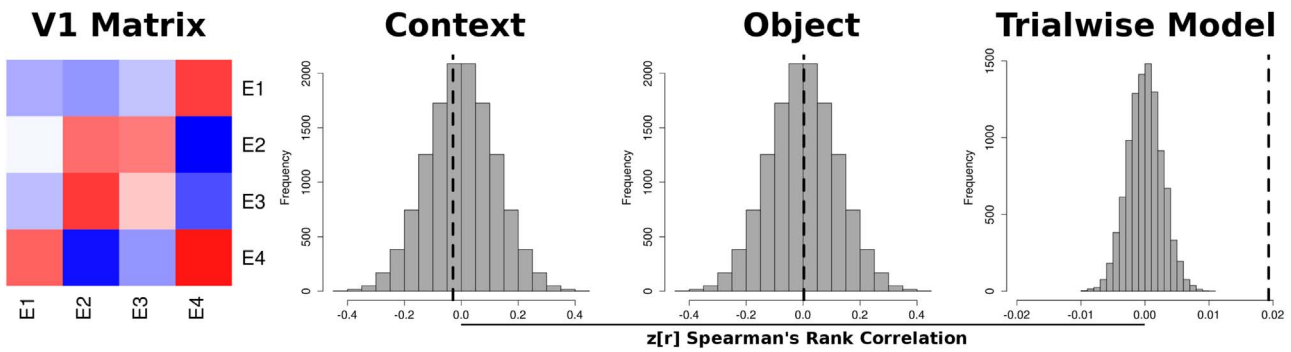

Fig. 5. The low-level confound was attenuated in Experiment 2. The V1 similarity matrix showed no sign of a relationship to either the context matrix (M=−0.029, p=0.81; BF_{01}=4.16) or the object matrix (M=0.0025, p=0.99; BF_{01}=4.30). Trialwise analysis revealed a significant relationship between the empirical V1 similarity matrix and the model V1 similarity matrix (M=0.019, p < 0.0001). The V1 similarity matrix was also related to the correct-response matrix (M=0.38, p < 0.005; data not shown).

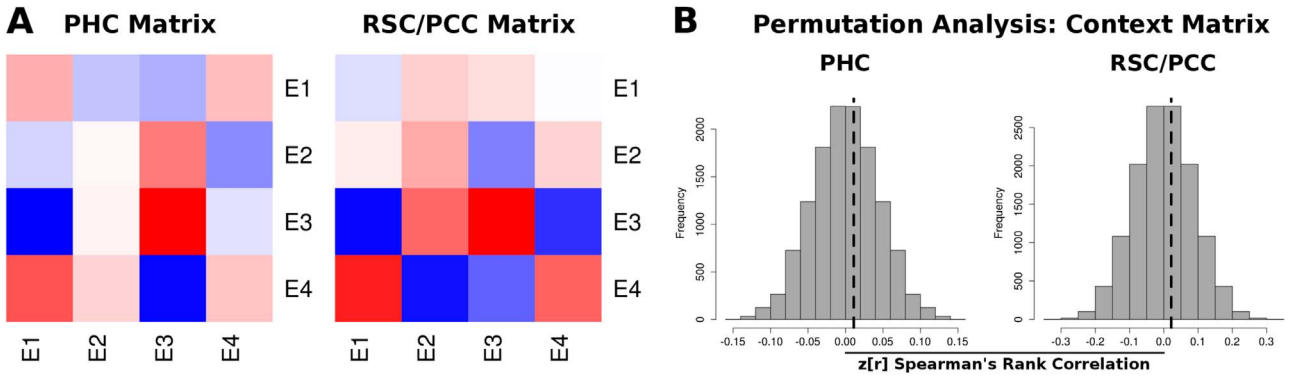

Fig. 6. No evidence for invariant context representation in PHC and RSC/PCC. The relationship between the context matrix and the PHC and the RSC/PCC similarity matrices failed to reach significance (PHC: M=0.011, p=0.80; BF_{01}=4.29; RSC/PCC: M=0.022, p=0.81; BF_{01}=4.23), suggesting that the relationships to the context matrix that were observed in Experiment 1 were highly influenced by the low-level sensory differences between the distinct contexts.

Fig. 5). We also performed a Bayes factor analysis to assess the evidence for the null hypothesis, where BF_{01} denotes the odds ratio of the null hypothesis relative to the alternative hypothesis (Rouder et al., 2009). We found that the relationship between the V1 similarity matrix and both the context matrix (BF_{01}=4.16) and the object matrix (BF_{01}=4.30) favored the null hypothesis, suggesting that the null hypothesis is approximately 4 times more likely than the alternative hypotheses. Moreover, a permutation analysis revealed a significantly weaker relationship between the V1 similarity matrix and the context matrix relative to Experiment 1 (p < 0.0001).

Next, a positive control analysis was performed to demonstrate a relationship between the empirical V1 similarity matrix and a model V1 similarity matrix (using HMAX; Theriault et al., 2011). We conducted a trialwise analysis, in which the V1 model was “shown” the same stimulus sequence as the participant. We generated trialwise similarity matrices for both the empirical data and the model, and we compared the empirical and model similarity matrices using Spearman's rank correlations. There was a significant relationship between the empirical and model V1 similarity matrices (M=0.019, p < 0.0001; Fig. 5). The relationship between the two matrices was numerically small but highly reliable, which was expected given that individual trial estimates of activity were used for the analysis. We also observed a significant relationship between the V1 similarity matrix and the correct-response matrix (M=0.38, p < 0.005), which was likely driven by the hemifield differences of the selected responses. Altogether, the results in V1 suggest that the low-level confound has been, at the very least, attenuated and the control analyses establish the quality of the data.

Investigation of the representation of context

We investigated whether the relationship to the context matrix maintained in our *a priori* ROIs. In contrast to Experiment 1, there was no sign of a relationship between the PHC and RSC/PCC similarity matrices and the context matrix (PHC: M=0.011, p=0.80; RSC/PCC: M=0.022, p=0.81; Fig. 6), and a within-hippocampus searchlight failed to reveal significant results. Additionally, the relationship between the PPA and the RS-Complex similarity matrices and the context matrix was severely diminished relative to Experiment 1, showing no sign of a relationship in PPA (M=0.029, p=0.69) and only a trend in RS-Complex (M=0.12, uncorrected p=0.051; Supplementary Fig. 4). A Bayes factor analysis revealed evidence in favor of the null hypothesis in PHC (BF_{01}=4.29), RSC/PCC (BF_{01}=4.23), PPA (BF_{01}=3.87), and RS-Complex (BF_{01}=2.16). Moreover, a permutation analysis revealed significantly weaker relationships between the context matrix and the similarity matrices in RSC/PCC, PPA, and RS-Complex relative to Experiment 1 (all p's < 0.0001; the decrease in PHC failed to reach significance: p=0.23). Thus, reducing (or eliminating) the low-level sensory differences across contexts markedly decreased the substantial contextual effects that we observed in Experiment 1.

Activation analysis: Events versus perceptual baseline

To ensure that the data were reliable and that the task was activating our ROIs, a standard activation analysis (events versus perceptual baseline) was conducted. This revealed significantly greater blood-oxygen-level dependent (BOLD) activity for events than the perceptual baseline task in PHC, RSC/PCC, PPA, and RS-Complex (PHC: t_9=6.4443, p < 0.0005; RSC/PCC: t_9=2.9574, p < 0.02; Left anatomical PPA: t_9=7.3008, p < 0.0001; Right anatomical PPA: t_9=6.6446, p < 0.0001; Left anatomical retrosplenial complex: t_9=4.5301, p < 0.005; Right anatomical retrosplenial complex: t_9=4.2866, p < 0.005). These findings suggest that our ROIs responded to the events relative to the baseline task even though there was little evidence for a relationship to the context matrix.

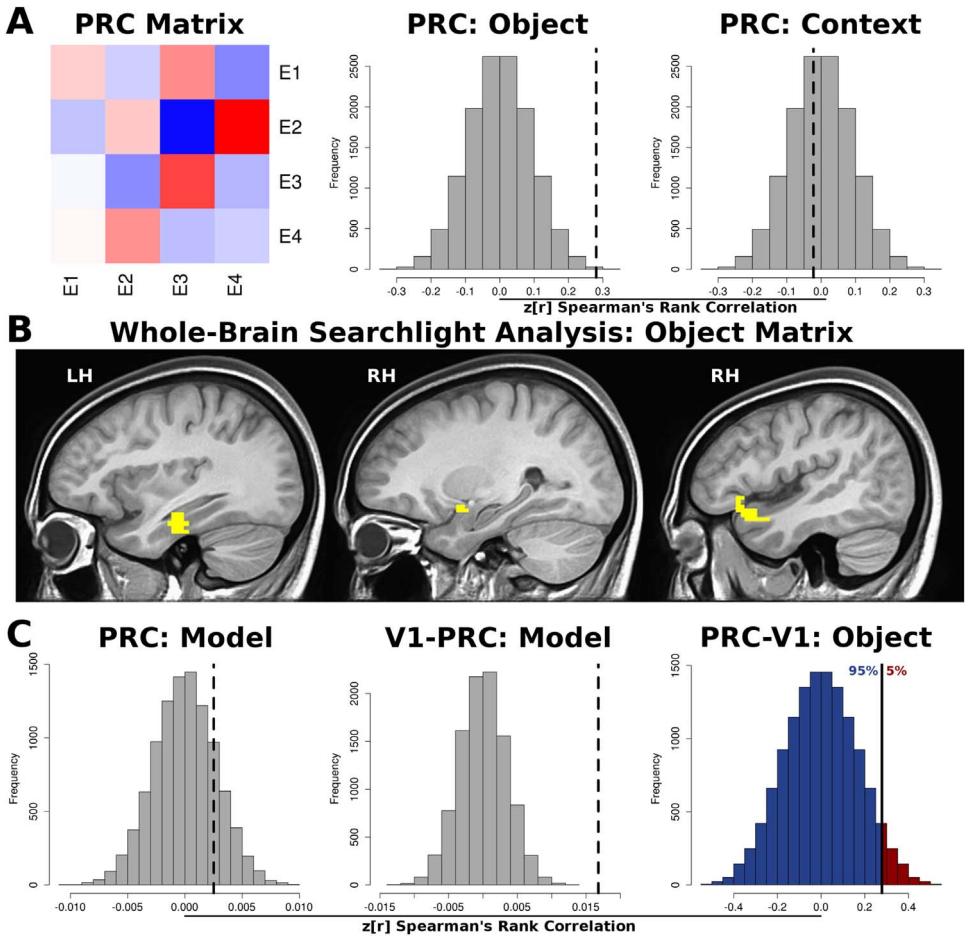

Fig. 7. Evidence for invariant object representation in perirhinal cortex. A) Permutation analysis revealed a significant relationship between the PRC similarity matrix and the object matrix (M=0.28, p=0.002) but the relationship to the context matrix failed to reach significance (M=−0.022, p=0.81). B) A whole-brain searchlight analysis revealed a cluster in the left PRC as well as two other clusters in the right anterior temporal lobe. C) The relationship between the trialwise PRC similarity matrix and the trialwise V1 model similarity matrix failed to reach significance (M=0.0025, p=0.36; left panel), the relationship between the trialwise V1 similarity matrix and the trialwise V1 model similarity matrix was significantly stronger that the relationship between the trialwise PRC similarity matrix and the trialwise V1 model similarity matrix (M=0.017, $p < 0.0001$; middle panel), and there was evidence that the relationship between the PRC similarity matrix and the object matrix was stronger than the relationship between the V1 similarity matrix and the object matrix (M=0.28, one-tailed $p < 0.05$; right panel). Similarly, the relationship between the left PRC cluster (left panel of B) similarity matrix and the object matrix was stronger than the relationship between the V1 similarity matrix and the object matrix (M=0.44, $p < 0.005$).

Investigation of object representation

To test the hypothesis that PRC carries object information (when the low-level stimulus features have been matched), we generated an object matrix that contains uniformly large values for events that share the same objects and uniformly small values for events that contain different objects (Fig. 2B). There was a significant relationship between the PRC similarity matrix and the object matrix (M=0.28, p=0.002), while the relationship to the context matrix failed to reach significance (M=−0.022, p=0.81; Fig. 7A). We next tested the prevalence of the relationship to the object matrix using a whole-brain searchlight analysis (3 voxel radius). A cluster was observed in left PRC, supporting the ROI-based results (57 voxel cluster, parametric voxel-wise $p < 0.01$; Fig. 7B). Additionally, two clusters were observed in the right anterior temporal cortex (53 and 42 voxels; middle and left panels of Fig. 7B, respectively). A follow-up analysis revealed that all three clusters themselves were significantly related to the object matrix (cluster 1: M=0.44, $p < 0.0005$; cluster 2: M=0.38, $p < 0.0005$; cluster 3: M=0.40, p=0.011, one participant was excluded from the cluster 3 analysis due to insufficient coverage). The follow-up analysis is circular but it is required to conclude that the cluster itself is informative (Etzel et al., 2013); therefore, the significant effects from this analysis bolster the conclusion that the clusters are related to the object matrix.

Testing for a double dissociation between PRC and V1

An important next step is to reveal a double dissociation between representations in PRC and V1. In contrast to V1, the relationship between the trialwise PRC and model V1 similarity matrices failed to reach significance (M=0.0025, p=0.36; left panel Fig. 7C). A nonparametric difference test revealed that the relationship between the empirical V1 similarity matrix and the model V1 similarity matrix was significantly stronger than the relationship between PRC similarity matrix and the model V1 similarity matrix (M=0.017, $p < 0.0001$; middle panel Fig. 7C). Finally, we observed some evidence that the relationship between the PRC similarity matrix and the object matrix was stronger than the relationship between the V1 similarity matrix and the object matrix (M=0.28, one-tailed $p < 0.05$; right panel Fig. 7C). Similarly, the relationship between the left PRC cluster (left panel of Fig. 7B) similarity matrix and the object matrix was stronger than the relationship between the V1 similarity matrix and the object matrix (M=0.44, $p < 0.005$). These results mitigate the possibility that object representations were inherited from V1.

Discussion

We conducted two experiments to investigate whether RSC/PCC and subregions of the MTL carry information about context, items, order, and

their conjunctions. We built upon a well-designed approach used in the rodent (Rajji et al., 2006; Komorowski et al., 2009, 2013; Navawongse and Eichenbaum, 2013; Tort et al., 2013; McKenzie et al., 2014; Farovik et al., 2015; Keene et al., 2016) and observed many analogous patterns using fMRI in humans. However, the results of our experiments highlight the importance of controlling for low-level sensory differences between the experimental conditions—here, different contexts and objects. Moreover, our results raise interesting questions about how to distinguish memory-related representations from processing-related representations. We will expand on these issues below.

Investigation of the representation of distinct contexts and objects

In Experiment 1, we observed a relationship between the context matrix and the similarity matrices in PHC, RSC/PCC, and the left posterior hippocampus. There are many reasons to expect that resolving what kind of information is coded for in the hippocampus will be more challenging when using fMRI data than when using electrophysiological data. However, the largest differences in patterns of activity in the rodent dorsal hippocampus (the homolog of human posterior hippocampus; McKenzie et al., 2014) and MTL cortical regions (Keene et al., 2016) were observed in response to changes to the context; therefore, the observed relationship to the context matrix in the left posterior hippocampus, PHC, and RSC/PCC is consistent with the results in the rodent. Furthermore, the RSC/PCC similarity matrix was related to the item-in-context matrix, and we observed a relationship between context and item-in-context representations in RSC/PCC and performance on the associative memory task. A critical question arises, however, about whether these results are confounded by the low-level visual differences between the stimuli, which we addressed by investigating representations in primary visual cortex.

The V1 similarity matrix was related to the context matrix and the item-in-context matrix, suggesting that our results in Experiment 1 could have been influenced by low-level visual differences between the stimuli. The V1 similarity matrix was also related to the item-in-order-in-context matrix, suggesting that early processing areas can exhibit distinct patterns of activity in response to a reconfiguration of the same stimuli. Moreover, we observed a relationship between representations in V1 and behavioral performance. Altogether, the results in V1 highlight the difficulty of disentangling the role of a brain region in memory versus processing (see The difficulty of dissociating processing-related representations from memory-related representations), especially given that MTL regions, and the hippocampus in particular, receive inputs from all sensory modalities.

We performed a whole-brain searchlight analysis to investigate the prevalence of the relationship to each of our matrices. The strongest effects were observed in the occipital lobes, where many voxels were related to not only the context matrix but also the item-in-context matrix and the item-in-order-in-context matrix. In hindsight, the whole-brain searchlight results and the V1 results are not surprising given the visual nature of the events, but these findings could provide a challenge to the conclusion that context plays an organizing role for processing within the MTL (McKenzie et al., 2014, 2015; Farovik et al., 2015; Keene et al., 2016). In the rodent version of the context-guided object association task, the contexts differed in terms of visual cues (black versus white), tactile cues (sandpaper versus cloth), and spatial cues (side of the apparatus) and the items differed in terms of the digging media (visual and tactile cues) and the olfactory cues, which makes it is difficult to interpret the degree to which regions differed in their responses. More generally, whenever two fixed (and constantly present) stimuli are compared (e.g., our two time-lapse videos), there will necessarily be low-level sensory differences between stimuli; thus, moving to stimulus sets can be advantageous. Additionally, computational models can be used to select stimuli that are devoid of a low-level visual confound prior to running an experiment. We provide an example of this approach in Experiment 2.

Investigation of invariant context representation

In Experiment 2, we tested whether the results from Experiment 1 would maintain in the absence of low-level sensory differences between the stimuli. We used a combined approach of image manipulation and computational model testing to eliminate the presence of category information from the low-level visual features. There was no sign of a relationship between the V1 similarity matrix and either the context matrix or the object matrix, suggesting that we adequately reduced the low-level visual confound. Importantly, two control analyses showed that V1 contained task-relevant information. First, a trialwise analysis revealed a relationship between the empirical and model V1 similarity matrices. Second, there was a relationship between the V1 similarity matrix and the correct-response matrix, which likely reflects the hemifield differences based on the left versus right response box being displayed (see bottom of Fig. 2A). The relationship between the correct-response matrix and the V1 similarity matrix provides further evidence that the low-level confound was attenuated—i.e., these results suggest that the image information canceled out and that the only reliable visual information was the filling in of the box in either hemifield. The positive control analysis results in V1 establish data quality.

We next investigated whether the hippocampus, PHC, and RSC/PCC contain invariant representations of context. We found a severely diminished relationship between the similarity matrices in these regions and the context matrix relative to Experiment 1. While these results suggest that the context effects in Experiment 1 were highly dependent on low-level visual differences between contexts, we do not take these results to be evidence that these regions do not contain invariant representations of context. The effect was clearly diminished (to the point of being undetectable), but we cannot say that it was entirely eliminated or that it could not be observed under other circumstances. We also acknowledge that the manipulation in Experiment 2 might have caused not only differences in the low-level sensory features between the contexts but could also have caused the memory content to be degraded relative to Experiment 1. Critically, participants performed well above chance on the associative memory task in Experiment 2, suggesting that they formed stable memory representations. Accordingly, it is clear that the substantial reduction in observed contextual information weakens the inference one might make from Experiment 1 or similar experiments that these regions contain generalized contextual information.

Future studies using electrophysiology and fMRI will be useful to help elucidate the conditions under which these regions exhibit invariant context coding. Additionally, designs that manipulate stimulus familiarity will be useful for understanding the involvement of subregions of the MTL and of RSC/PCC in context representation as well as the role of memory in the formation of these representations. For example, it is possible that our participants did not have real-world experience with Saint Peter's Basilica or with the U.S. Capitol Building, which led to little evidence for invariant context representation (for evidence for the importance of landmark familiarity in PPA see: Marchette et al., 2015). Importantly, all of our scene ROIs showed greater BOLD activation for events compared to the perceptual baseline task, suggesting that they were activated by the associative memory task.

Investigation of invariant object representation

In contrast to the relative lack of findings for invariant context representation, we observed a relationship between the PRC similarity matrix and the object matrix in Experiment 2. These results support the hypothesis that PRC is involved in the representation of objects (Burwell, 2000; Davachi, 2006; Knierim et al., 2006; Diana et al., 2007; Eichenbaum et al., 2007; Wixted and Squire, 2011; Ranganath and Ritchey, 2012). A whole-brain searchlight analysis revealed only three clusters, the first of which was centered in left PRC and the other two of which were located in right anterior temporal lobe (one of which was in

close proximity to PRC, containing some overlapping voxels with the temporal-polar portion of PRC). Importantly, a follow-up analysis revealed that the clusters themselves were related to the object matrix. These results corroborate the ROI-based approach and suggest that the relationship to the object matrix is relatively exclusive to the more anterior portions of the MTL.

We propose that an important next step is to confirm that there is a dissociation between representations in PRC and early visual areas, which mitigates the possibility that representations were inherited from earlier processing regions. Accordingly, we used nonparametric difference tests to investigate whether there was a double dissociation between representations in PRC and V1. After establishing that PRC was not related to the trialwise model V1 similarity matrix, we showed that the relationship between the empirical V1 similarity matrix and the model V1 similarity matrix was stronger than the relationship between the PRC similarity matrix and the model V1 similarity matrix. Next, we provided evidence to suggest that the relationship between the PRC similarity matrix and the object matrix was stronger than the relationship between the V1 similarity matrix and the object matrix. Taken together, these results suggest that PRC contains representations of car keys and house keys that generalize across multiple exemplars and these representations are not inherited from V1. Future studies can investigate whether these effects depend on the associative memory task or would be observed more generally.

An alternative to our difference tests would be to use a partial correlation approach. For example, Clarke and Tyler (2014) used a partial correlation representational similarity analysis to investigate representations throughout the visual stream. One benefit of this approach is that it is straightforward to implement in experiments with larger and less constrained stimulus sets than we used in Experiment 2, similar to the stimulus set used by Clarke and Tyler (2014). Their results suggest that PRC represents fine-grained semantic information about individual objects. Importantly, this effect maintained while holding the effect of a model of V1 representations constant. Furthermore, they used a modeling approach to show that BOLD activation in PRC was modulated by the confusability of objects. Taken together, their results suggest a role for PRC in fine-grained semantic representations of objects (for review see: Clarke and Tyler, 2015). Similarly, research in patient populations has revealed a necessary role for the anterior temporal cortex, and PRC in particular, in naming highly confusable objects (Kivisaari et al., 2012; Wright et al., 2015). Our results extend previous findings by showing invariant representation of subordinate object categories (car keys versus house keys). The differences in experimental design between our Experiment 2 and the approach used by Clarke and Tyler (2014) provide converging evidence for the representation of fine-grained category information in PRC. These experimental designs can reveal complementary information, thus we propose that both approaches will be useful for future studies.

High-level cognitive representations versus low-level sensory representations

While many neuroimaging studies have investigated the involvement of high-level cortical regions in the processing of categories of visual stimuli, relatively few fMRI studies have investigated whether low-level stimulus properties modulate the responses in these regions. Recent studies have found evidence for retinotopic responses in PPA (Arcaro et al., 2009; Silson et al., 2015). Specifically, these studies reported that PPA responded preferentially to stimuli that were presented in the upper contralateral visual field (relative to other locations). PPA has also been shown to exhibit greater activation in response to images with high spatial frequencies compared to low spatial frequencies (Rajimehr et al., 2011) and in response to cardinal orientations (i.e., vertical or horizontal) relative to oblique orientations (Nasr and Tootell, 2012). Moreover, Nasr et al. (2014) showed that scene stimuli from previously published studies had larger values of rectilinearity (i.e., energy at 90 degree angles) than stimuli from other categories, and they reported greater PPA activation in response to rectilinear stimuli compared to round stimuli. Lescroart et al. (2015) showed that a Fourier power model, which measures the spectral features of images, accounted for a large amount of variance in PPA and RS-Complex. Additionally, recent studies have found that the representational similarity matrices in category-selective cortical regions were related to the low-level features of the stimuli (as measured using the GIST model) for object categories (Rice et al., 2014; Watson et al., 2016b) and scene categories (Watson et al., 2014; for review see Andrews et al., 2015). Finally, image filtering (high-pass versus low-pass) of scene images has been shown to influence the representational similarity matrices in scene-selective regions (Watson et al., 2016a).

Recent studies have suggested that hippocampal place cells might also be influenced by low-level sensory cues. Geva-Sagiv et al. (2016) recorded from the bat hippocampus (CA1 and subiculum) and provided evidence that place cell maps between vision- and echolocation-based navigation of the same environment are unrelated, raising questions about whether place cells represent the neural instantiation of an invariant spatial map or whether they contain sensory-specific signals. Additionally, studies in the rodent that have recorded simultaneously from V1 and the hippocampus have shown “place cell” activity in V1 (Ji and Wilson, 2007), which was also shown to precede place cell activity in the hippocampus (Haggerty and Ji, 2015). Thus, a more complete understanding of the sensory versus cognitive influences on representations in the MTL can be realized by studies that simultaneously investigate representations in sensory cortices and the MTL (also see Investigation of invariant object representation). Taken together, our results and the results of the studies discussed here suggest that we need to be very careful in designing experiments—especially visual experiments in humans—to ensure that high-level cognitive representations are not confounded by low-level sensory features (for similar arguments regarding the importance of controlling for confounds in multivariate pattern analysis see: Davis and Poldrack, 2013; Todd et al., 2013).

In our daily lives we can use low-level sensory features to provide us with information about our current context. However, in experiments that attempt to decode the representation of context, such features become confounds. We propose that there are two classes of experimental designs that allow the investigation of the representation of context while minimizing the presence of low-level confounds: 1) investigations of stable representations across different instances of the same context (the aggregate of which are not discriminable based on the low-level features), 2) investigations of distinct representations while the animal is exposed to the same sensory cues but with different contextual cues (which are not concurrently available to the animal's sensory system). We used the first approach in Experiment 2 (also see: Marchette et al., 2015) because it provided a straightforward method for controlling the low-level features of both the contexts and the objects; however, other studies have used the second approach to investigate representations in the hippocampus.

Previous studies have shown differential place cell firing in the rat hippocampus in response to the same physical stimuli as a result of behavioral conditions (Wood et al., 2000; Smith and Mizumori, 2006b) and expectations (Skaggs and McNaughton, 1998), suggesting that internal context can play a dramatic role even within the same physical environment (for reviews see: Smith and Mizumori, 2006a; Smith and Bulkin, 2014). Additionally, recent studies in humans (Hsieh et al., 2014) and rats (Allen et al., 2016) have shown that the hippocampus carries conjunctive item-in-sequence information (i.e., temporal context). Collectively, the experimental designs of these studies are interesting because they allow the investigation of the influence of context without manipulating the physical stimuli (i.e., the low-level cues are controlled because they are identical across contexts). These examples suggest that context undoubtedly plays an important role in the organization of

representations in the hippocampus, however future experiments should test the degree to which sensory versus cognitive factors influence representations in both rats and humans. For example, representations of sinusoidal gratings have been shown to be modulated by attention in the lateral geniculate nucleus (i.e., the first synapse from the retina) during a fixation task (Ling et al., 2015). These results demonstrate the potential for low-level sensory processing to change as a result of attention and memory. Therefore, it is possible that confounds similar to those observed in Experiment 1 can persist even in experiments that expose animals to the same physical stimuli.

We acknowledge that the task used in the present experiments is not a traditional single-shot episodic memory task. However, we argue that our results can be brought to bear to raise similar concerns regarding episodic memory tasks. For example, one approach for turning Experiment 1 into a single-shot episodic memory task would be to present a sequence of objects (each of which would be shown only once) embedded within a distinct context (e.g., a room, a time-lapse video of a location [as in Experiment 1]). Then, participants would be exposed to a second sequence of images embedded within a second distinct context. Participants would then perform a retrieval task in which they are asked to decide the context in which each image was presented. However, the same low-level confound would be present because such a study would use distinct stimuli for the contexts. In fact, in the limit in which the representations of two distinct experiences are examined, a low-level sensory confound will be guaranteed to be present—i.e., two experiences that differ in terms of "what-where-when" will always differ in both the high-level cognitive features and the low-level sensory features. Accordingly, we suggest that the experimental design of such an experiment should implement an approach similar to that taken in Experiment 2—i.e., match the low-level features between the stimuli that comprise the events in addition to using large stimulus sets. Therefore, while the task used in the present experiments is not an episodic memory task per se, we argue that our results have implications for memory experiments writ large and are not limited to any one type of memory task.

The difficulty of dissociating processing-related representations from memory-related representations

Separate from the difficulty of dissociating low-level sensory representations and high-level cognitive representations, there is an additional difficulty in dissociating processing-related representations (which are comprised of both low- and high-level representations of stimulus-related information) from memory-related representations. For example, the relationship that we observed in Experiment 1 between representations in V1 and behavioral performance raises interesting questions about how to infer whether brain-behavior relationships are indicative of memory or processing; i.e., while V1 clearly *could* play a role in memory, it is not traditionally included in models of declarative memory. Similar challenges have been raised from the results an electrophysiological study in which monkeys were trained to learn an association between static cue stimuli (up and down arrows) and the direction of movement of stimuli (upward moving dots and downward moving dots; the up/down arrow and movement pairing was counterbalanced across monkeys; Schlack and Albright, 2007). They found that area MT, which is known to show motion-selective responses, exhibited selective responses to the cue stimuli after learning. Specifically, neurons began responding to the cue stimuli in a similar manner to their responses to the associated motion stimuli. The authors suggested that the selective responses to the cue could reflect either the dynamic reorganization of area MT (i.e., neurons in MT were modified to represent non-motion stimuli) or "recall-related" activity (i.e., motion-related activity was reactivated as a result of associative memory retrieval in another brain region), concluding that the latter option is more likely (also see: Albright, 2012).

Findings from human neuroimaging studies have also demonstrated the difficulty of discriminating between memory and processing. For example, Lee et al. (2012) and Naselaris et al. (2015) trained participants to vividly imagine images. Lee et al. (2012) found that classifiers that were trained on patterns of activity during perception were able to decode the identity of objects during imagery in early visual areas, and decoding accuracy correlated with the vividness of imagery. Moreover, Naselaris et al. (2015) found that an encoding model that was based solely on the low-level features of the images (the model used Gabor wavelets; for review see: Naselaris et al., 2011) could be used to successfully decode the identity of images during mental imagery (i.e., retrieval of those stimuli). While virtually all memory models posit that recollection would result in the reactivation of the cortical regions that were involved during encoding (and presumably the low-level features as well), the results of these studies reveal the difficulty of inferring whether brain-behavior relationships are indicative of memory that is localized to the ROI or whether such relationships reflect perceptual reactivation from other regions that encoded or retrieved the memory. Moreover, separate from the difficulty of dissociating processing-related representations from memory-related representations, these studies suggest that memory retrieval paradigms are not immune to the challenges associated with discriminating between low-level versus high-level representations (see High-level cognitive representations versus low-level sensory representations) because memory retrieval can reactivate the low-level confound in early visual areas.

The studies discussed thus far are inherently correlational; however, similar challenges exist for more causal approaches. For example, optogenetic techniques allow researchers to "tag" neurons that are active during encoding and then to later selectively silence or reactivate those cells. Tanaka et al. (2014) tagged hippocampal cells that were active during contextual fear conditioning. They later silenced the same hippocampal cells during re-exposure to the previously shocked context, which was sufficient to decrease both freezing (memory-related behavior) and the reactivation of previously active cortical cells. Cowansage et al. (2014) used similar techniques to tag active neurons in retrospenial cortex during contextual fear conditioning. They demonstrated that pharmacological inactivation of the dorsal hippocampus decreased freezing upon re-exposure to the previously shocked context. Interestingly, selective reactivation of tagged cells in retrosplenial cortex during pharmacological inactivation of the hippocampus was sufficient to elicit a similar amount of freezing to a control group with an intact hippocampus. Accordingly, these studies provide causal evidence that specific neuronal ensembles in the hippocampus and retrosplenial cortex are involved in processing contextual information (or at least the low-level features that comprise the context); however, even these studies do not provide evidence that these ensembles reflect memory representations (Mayford and Reijmers, 2016). In fact, Mayford and Reijmers (2016, pg. 14) provided a convincing example of the difficulty of discriminating between memory and processing in these studies, stating: "[I]f one could stimulate retinal neurons in precisely the same pattern as when an animal explored a specific context, then the animal would presumably perceive that they were in the box, and freeze if fear conditioned, but we would not suggest that the engram for that memory lies in the retina."

Taken together, changes from a pre-experimental state (e.g., Schlack and Albright, 2007), retrieval-based reactivation (including stimulus-specific reactivation; e.g., Lee et al., 2012; Naselaris et al., 2015), brain-behavior relationships (e.g., Experiment 1; Lee et al., 2012), and changes in memory-related behavior during optogenetic silencing or reactivation (e.g., Tanaka et al., 2014; Cowansage et al., 2014) are not sufficient to conclude that decoded representations are indicative of memory. Accordingly, invasive studies can address whether the results of recording (e.g., electrophysiology, fMRI) and stimulation (e.g., optogenetic reactivation) studies reflect processing or memory. For example, Mayford and Reijmers (2016) suggested that the combination of optogenetic silencing and reactivation with techniques for studying learning-induced synaptic reorganization will be

important for understanding how neuronal ensembles relate to memory versus processing.

Acknowledgments

This work was supported by NIA grant R01 AG034613 and NIH grant R01 MH085828 awarded to Craig Stark. We thank Patricia Place and Samantha Rutledge for assistance with data collection. We thank Shauna Stark, Veronique Boucquey, John Guzowski, Zach Reagh, and Jared Stokes for helpful discussions. We thank Norbert Fortin and three anonymous reviewers for helpful comments on earlier versions of this manuscript.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.neuroimage.2017.04.019.

References

Aggleton, J.P., 2010. Understanding retrosplenial amnesia: insights from animal studies. Neuropsychologia 48 (8), 2328–2338.
Aguirre, G.K., 2007. Continuous carry-over designs for fMRI. Neuroimage 35 (4), 1480–1494.
Aguirre, G.K., Mattar, M.G., Magis-Weinberg, L., 2011. de Bruijn cycles for neural decoding. Neuroimage 56 (3), 1293–1300.
Albright, T.D., 2012. On the perception of probable things: neural substrates of associative memory, imagery, and perception. Neuron 74 (2), 227–245.
Alexander, A.S., Nitz, D.A., 2015. Retrosplenial cortex maps the conjunction of internal and external spaces. Nat. Neurosci. 18 (8), 1143–1151.
Allen, T.A., Salz, D.M., McKenzie, S., Fortin, N.J., 2016. Nonspatial sequence coding in CA1 neurons. J. Neurosci. 36 (5), 1547–1563.
Andrews, T.J., Watson, D.M., Rice, G.E., Hartley, T., 2015. Low-level properties of natural images predict topographic patterns of the neural response in the ventral visual pathway. J. Vision. 15 (7), 3.
Arcaro, M.J., McMains, S.A., Singer, B.D., Kastner, S., 2009. Retinotopic organization of human ventral visual cortex. J. Neurosci. 29 (34), 10638–10652.
Auger, S.D., Maguire, E.A., 2013. Assessing the mechanism of response in the retrosplenial cortex of good and poor navigators. Cortex 49 (10), 2904–2913.
Auger, S.D., Zeidman, P., Maguire, E.A., 2015. A central role for retrosplenial cortex in de novo enviromental learning. eLife 4, e09031.
Avants, B., Duda, J.T., Kim, J., Zhang, H.Y., Pluta, J., Gee, J.C., Whyte, J., 2008. Multivariate analysis of structural and diffusion imaging in traumatic brain injury. Acad. Radiol. 15, 1360–1375.
Bar, M., Aminoff, E., 2003. Cortical analysis of visual context. Neuron 38 (2), 347–358.
Brainard, D.H., 1997. The psychophysics toolbox. Spat. Vision. 10, 443–446.
Bucci, D.J., Robinson, S., 2014. Toward a conceptualization of retrohippocampal contributions to learning and memory. Neurobiol. Learn. Mem. 116, 197–207.
Bulkin, D.A., Law, M., Smith, D.M., 2016. Placing memories in context: hippocampal representations promote retrieval of appropriate memories. Hippocampus. http://dx.doi.org/10.1002/hipo.22579.
Burwell, R.D., 2000. The parahippocampal region: corticocortical connectivity. Ann. NY. Acad. Sci. 911, 25–42.
Chen, L.L., Lin, L.-H., Green, E.J., Barnes, C.A., McNaughton, B.L., 1994. Head-direction cells in the rat posterior cortex. i. anatomical and behavioral modulation. Exp. Brain Res. 101, 8–23.
Chen, Y., Namburi, P., Elliot, L.T., Heinzie, J., Soon, C.S., Chee, M.W., Haynes, J.D., 2011. Cortical surface-based searchlight decoding. Neuroimage 56 (2), 582–592.
Cho, J., Sharp, P.E., 2001. Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. Behav. Neurosci. 115 (1), 3–25.
Clarke, A., Tyler, L.K., 2014. Object-specific semantic coding in human perirhinal cortex. J. Neurosci. 34 (14), 4766–4775.
Clarke, A., Tyler, L.K., 2015. Understanding what we see: how we derive meaning from vision. Trends Cogn. Sci. 19 (11), 677–687.
Cohen, N.J., Ryan, J., Hunt, C., Romine, L., Wszalek, T., Nash, C., 1999. Hippocampal system and declarative (relational) memory: summarizing the data from functional neuroimaging studies. Hippocampus 9 (1), 83–98.
Coutanche, M.N., Thompson-Schill, S.L., 2012. The advantage of brief fMRI acquisition runs for multi-voxel pattern detection across runs. Neuroimage 61 (4), 1113–1119.
Coutanche, M.N., Thompson-Schill, S.L., 2013. Informational connectivity: identifying synchronized discriminability of multi-voxel patterns across the brain. Front. Hum. Neurosci. 7 (15), 1–14.
Cowansage, K.K., Shuman, T., Dillingham, B.C., Chang, A., Colshani, P., Mayford, M., 2014. Direct reactivation of a coherent neocortical memory of context. Neuron 84 (2), 432–441.
Cox, R.W., 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Comput. Biomed. Res. 29 (3), 162–173.
Davachi, L., 2006. Item, context and relational episodic encoding in humans. Curr. Opin. Neurobiol. 16, 693–700.
Davis, T., Poldrack, R.A., 2013. Measuring neural representations with fMRI: practices and pitfalls. Ann. NY Acad. Sci. 1296, 108–134.
Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., Albert, M.S., Killiany, R.J., 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. Neuroimage 31, 968–980.
Destrieux, C., Fischl, B., Dale, A., Halgren, E., 2010. Automatic parcellation of human cortical gyri and sulci using standard anatomical nomenclature. Neuroimage 53 (1), 1–15.
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 (9), 379–386.
Eichenbaum, H., Yonelinas, A.P., Ranganath, C., 2007. The medial temporal lobe and recognition memory. Annu. Rev. Neurosci. 30, 123–152.
Ennaceur, A., Neave, N., Aggleton, J.P., 1997. Spontaneous object recognition and object location memory in rats: the effects of lesions in the cingulate cortices, the medial prefrontal cortex, the cingulum bundle and the fornix. Exp. Brain Res. 113, 509–519.
Epstein, R., Kanwisher, N., 1998. A cortical representation of the local visual environment. Nature 392, 598–601.
Epstein, R.A., Parker, W.E., Feiler, A.M., 2007. Where am I now? Distinct roles for parahippocampal and retrosplenial cortices in place recognition. J. Neurosci. 27 (23), 6141–6149.
Ernst, M.D., 2004. Permutation methods: a basis for exact inference. Stat. Sci. 19, 676–685.
Etzel, J.A., Zacks, J.M., Braver, T.S., 2013. Searchlight analysis: promise, pitfalls, and potential. NeuroImage 78, 261–269.
Etzel, J.A., 2015. MVPA permutation schemes: permutation testing for the group level. Pattern Recognition in NeuroImaging (PRNI), 2015 International Workshop on, pp. 65–68.
Farovik, A., Place, R.J., McKenzie, S., Porter, B., Munro, C.E., Eichenbaum, H., 2015. Orbitofrontal cortex encodes memories within value-based schemas and represents contexts that guide memory retrieval. J. Neurosci. 35 (21), 8333–8344.
Geva-Sagiv, M., Romani, S., Las, L., Ulanovsky, N., 2016. Hippocampal global remapping for different sensory modalities in flying bats. Nat. Neurosci. 19 (7), 952–958.
Griffin, G., Holub, A., Perona, P., 2007. Caltech-256 object category dataset. Technical Report 7694, California Institute of Technology.
Haggerty, D.C., Ji, D., 2015. Activities of visual cortical and hippocampal neurons co-fluctuate in freely moving rats during spatial behavior. eLife. http://dx.doi.org/10.7554/eLife.08902.
Hanke, M., Halchenko, Y.O., 2012. Neuroscience runs on GNU/Linux. Neuroinformatics, 5–8.
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 (1), 37–53.
Haxby, J.V., Gobbini, M.I., Lurey, M.L., Ishai, A., Schouten, J.L., P, P., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science 293, 2425–2430.
Hinds, O.P., Rajendran, N., Polimeni, J.R., Augustinack, J.C., Wiggins, G., Wald, L.L., Rosas, H.D., Potthast, A., Schwarts, E.L., Fischl, B., 2008. Accurate prediction of V1 location from cortical folds in a surface coordinate system. Neuroimage 39 (4), 1585–1599.
Hsieh, L.T., Gruber, M.J., Jenkins, L.J., Ranganath, C., 2014. Hippocampal activity patterns carry information about objects in temporal context. Neuron 81 (5), 1165–1178.
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 24 (11), 1394–1403.
Ji, D., Wilson, M.A., 2007. Coordinated memory replay in the visual cortex and hippocampus during sleep. Nat. Neurosci. 10 (1), 100–107.
Julian, J.B., Fedorenko, E., Webster, J., Kanwisher, N., 2012. An algorithmic method for functionally defining regions of interest in the ventral visual pathway. Neuroimage 60 (4), 2357–2364.
Keene, C.S., Bladon, J., McKenzie, S., Liu, C.D., O'Keefe, J., Eichenbaum, H., 2016. Complimentary functional organization of neuronal activity patterns in the perirhinal, lateral entorhinal, and medial entorhinal cortices. J. Neurosci. 36 (13), 3660–3675.
Kivisaari, S.L., Tyler, L.K., Monsch, A.U., Taylor, K.I., 2012. Medial perirhinal cortex disambiguates confusable objects. Brain 135 (Pt 12), 3757–3769.
Knierim, J.J., Lee, I., Hargreaves, E.L., 2006. Hippocampal place cells: parallel input streams, subregional processing, and implications for episodic memory. Hippocampus 16, 755–764.
Komorowski, R.W., Manns, J.R., Eichenbaum, H., 2009. Robust conjunctive item-place coding by hippocampal neurons parallels what happens where. J. Neurosci. 29 (31), 9918–9928.
Komorowski, R.W., Garcia, C.G., Wilson, A., Hattori, S., Howard, M.W., Eichenbaum, H., 2013. Ventral hippocampal neurons are shaped by experience to represent behaviorally relevant contexts. J. Neurosci. 33 (18), 8079–8087.
Kriegeskorte, N., Mur, M., Bandettini, P., 2008a. Representational similarity analysis – connecting the branches of systems neuroscience. Front. Syst. Neurosci. 2 (4).
Kriegeskorte, N., Mur, M., Ruff, D.A., Kiani, R., Bodurka, J., Esteky, H., Tanaka, K., Bandettini, P.A., 2008b. Matching categorical object representations in inferior temporal cortex of man and monkey. Neuron 60 (6), 1126–1141.
Krzanowski, W.J., 2000. Principles of multivariate analysis: a user's perspective. Oxford University Press.
Lavenex, P., Amaral, D.G., 2000. Hippocampal-neocortical interaction: a hierarchy of associativity. Hippocampus 10, 420–430.
Law, J.R., Flanery, M.A., Wirth, S.W., Yanike, M., Smith, A.C., Frank, L.M., Suzuki, W.A., Brown, E.N., Stark, C.E.L., 2005. Functional magnetic resonance imaging activity during the gradual acquisition and expression of paired-associate memory. J. Neurosci. 25 (24), 5720–5729.
Lazebnik, S., Schmid, C., Ponce, J., 2006. Beyond bags of features: spatial pyramid matching for recognizing natural scene categories. In: Proceedings of the 2006 IEEE Computer Science Society Conference on Computer Vision and Pattern Recognition (CVPR'06), 2. pp. 2169–2178.

Lee, S.-H., Kravitz, D.J., Baker, C.I., 2012. Disentangling visual imagery and perception of real-world objects. Neuroimage 59 (4), 4064–4073.
Lescroart, M.D., Stansbury, D.E., Gallant, J.L., 2015. Fourier power, subjective distance, and object categories all provide plausible models of BOLD responses in scene-selective areas. Front. Comput. Neurosci. 9, 135.
Liang, M., Mouraux, A., Hu, L., Iannetti, G.D., 2013. Primary sensory cortices contain distinguishable spatial patterns of activity for each sense. Nat. Commun. 4, 1979.
Ling, S., Pratte, M.S., Tong, F., 2015. Attention alters orientation processing in the human lateral geniculate nucleus. Nat. Neurosci. 18 (4), 496–498.
Marchette, S.A., Vass, L.K., Ryan, J., Epstein, R.A., 2014. Anchoring the neural compass: coding of local spatial reference frames in human medial parietal lobe. Nat. Neurosci. 17 (11), 1580–1606.
Marchette, S.A., Vass, L.K., Ryan, J., Epstein, R.A., 2015. Outside looking in: landmark generalization in the human navigational system. J. Neurosci. 35 (44), 14896–14908.
Mayford, M., Reijmers, L., 2016. Exploring memory representations with activity-based genetics. Cold Spring Harb. Perspect. Biol. 8, a021832.
McKenzie, S., Frank, A.J., Kinsky, N.R., Porter, B., Rivière, P.D., 2014. Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. Neuron 83 (1), 202–215.
McKenzie, S., Keene, C.S., Farovik, A., Bladon, J., Place, R., Komorowski, R., Eichenbaum, H., 2016. Representation of memories in the cortical-hippocampal system: results from the application of population similarity analyses. Neurobiol. Learn. Mem. 134 (Pt A), 178–191.
Mishkin, M., Suzuki, W.A., Gadian, D.G., Vargha-Khadem, F., 1997. Hierarchical organization of cognitive memory. Philisophical Trans. R. Soc. B 352, 1461–1467.
Morris, R., 2006. Theories of hippocampal function. In: Anderson, P., Morris, R., Amaral, D., Bliss, T., O'Keefe, J. (Eds.), The Hippocampus Book1st ed.. Oxford University Press, New York, 581–714.
Nadel, L., Willner, J., 1980. Context and conditioning: a place for space. Physiol. Psychol. 8 (2), 218–228.
Naselaris, T., Kay, K.N., Nishimoto, S., Gallant, J.L., 2011. Encoding and decoding in fMRI. Neuroimage 56, 400–410.
Naselaris, T., Olman, C., Stansbury, D., Ugurbil, K., Gallant, J., 2015. A voxel-wise encoding model for early visual cortex decodes mental images of remembered scenes. Neuroimage 105, 215–228.
Nasr, S., Tootell, R.B.H., 2012. A cardinal orientation bias in scene-selective visual cortex. J. Neurosci. 32 (43), 14921–14926.
Nasr, S., Echavarria, C.E., Tootell, R.B.H., 2014. Thinking outside the box: rectilinear shapes selectively activate scene-selective cortex. J. Neurosci. 34 (20), 6721–6735.
Navawongse, R., Eichenbaum, H., 2013. Distinct pathways for rule-based retrieval and spatial mapping of memory representations in hippocampal neurons. J. Neurosci. 33 (3), 1002–1013.
Oliva, A., Torralba, A., 2001. Modeling the shape of the scene: a holistic representation of the spatial envelope. Int. J. Comput. Vision. 42 (3), 145–175.
Park, S., Chun, M.M., 2009. Different roles of the parahippocampal place area (PPA) and retrosplenial cortex (RSC) in panoramic scene perception. Neuroimage 47, 1747–1756.
Parron, C., Save, E., 2004. Comparison of the effects of entorhinal and retrosplenial cortical lesions on habituation, reaction to spatial and non-spatial changes during object exploration in the rat. Neurobiol. Learn. Mem. 82, 1–11.
Pelli, D.G., 1997. The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spat. Vision. 10, 437–442.
Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. Neuroimage 59 (3), 2142–2154.
Rajimehr, R., Devaney, K.J., Bilenko, N.Y., Young, J.C., Tootell, R.B.H., 2011. The "parahippocampal place area" responds preferentially to high spatial frequencies in humans and monkeys. PLoS Biol. 9 (4).
Rajji, T., Chapman, D., Eichenbaum, H., Greene, R., 2006. The role of CA3 hippocampal NMDA receptors in paired associate learning. J. Neurosci. 26 (3), 908–915.
Ranganath, C., Ritchey, M., 2012. Two cortical systems for memory-guided behaviour. Nat. Rev. Neurosci. 13, 713–726.
Rice, G.E., Watson, D.M., Hartley, T., Andrews, T.J., 2014. Low-level image properties of visual objects predict patterns of neural response across category-selective regions of the ventral visual pathway. J. Neurosci. 34 (26), 8837–8844.
Rissman, J., Gazzaley, A., D'Esposito, M., 2004. Measuring functional connectivity during distinct stages of a cognitive task. Neuroimage 23, 752–763.
Rouder, J.N., Speckman, P.L., Sun, D., Morey, R.D., Iverson, G., 2009. Bayesian t tests for accepting and rejecting the null hypothesis. Psychon. Bull. Rev. 16 (2), 225–237.
Saad, Z.S., Glen, D.R., Chen, G., Beauchamp, M.S., Desai, R., Cox, R.W., 2009. A new method for improving functional-to-structural alignment using local Pearson correlation. Neuroimage 44, 839–848.
Schlack, A., Albright, T.D., 2007. Remembering of motion: neural correlates of associative plasticity and motion recall in cortical area MT. Neuron 53 (6), 881–890.
Scoville, W.B., Milner, B., 1957. Loss of recent memory after bilateral hippocampal lesions. J. Neurol., Neurosurg. Psychiatry 20, 11–21.
Silson, E.H., Chan, A.W.Y., Reynolds, R.C., Kravitz, D.J., Baker, C.I., 2015. A retinotopic basis for the division of high-level scene processing between lateral and ventral human occipitotemporal cortex. J. Neurosci. 35 (34), 11921–11935.
Skaggs, W.E., McNaughton, B.L., 1998. Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. J. Neurosci. 18 (20), 8455–8466.
Smith, D.M., Mizumori, S.J.Y., 2006a. Hippocampal place cells, context, and episodic memory. Hippocampus 16 (9), 716–729.
Smith, D.M., Mizumori, S.J.Y., 2006b. Learning-related development of context-specific neuronal responses to places and events: the hippocampal role in context processing. J. Neurosci. 26 (12), 3154–3163.
Smith, D.M., Bulkin, D.A., 2014. The form and function of hippocampal context representations. Neurosci. Biobehav. R. 40, 52–61.
Stelzer, J., Chen, Y., Turner, R., 2013. Statistical inference and multiple testing correction in classification-based multi-voxel pattern analysis (MVPA): random permutations and cluster size control. Neuroimage 65, 69–82.
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 (2), 347–354.
Theriault, C., Thome, N., Cord, M., 2011. HMAX-S: deep scale representation for biologically inspired image categorization. In Proc. IEEE Int. Conf. Image Process., Sep 2011, pp. 1261–1264.
Todd, M.T., Nystrom, L.E., Cohen, J.D., 2013. Confounds in multivariate pattern analysis: theory and rule representation case study. Neuroimage 77, 157–165.
Tort, A.B.L., Komorowski, R., Kopell, N., Eichenbaum, H., 2013. A mechanism for the formation of hippocampal neuronal firing patterns that represent what happens where. Learn. Mem. 18 (11), 718–727.
Turner, B.O., Mumford, J.A., Poldrack, R.A., Ashby, F.G., 2012. Spatiotemporal activity estimation for multivoxel pattern analysis with rapid event-related designs. Neuroimage 62 (3), 1429–1438.
Valenstein, E., Bowers, D., Verfaellie, M., Heilman, K.M., Day, A., Watson, R.T., 1987. Retrosplenial amnesia. Brain 110, 1631–1646.
Van Dijk, K.R.A., Sabuncu, M.R., Buckner, R.L., 2012. The influence of head motion on intrinsic functional connectivity. Neuroimage 59 (1), 431–438.
Vann, S.D., Aggleton, J.P., 2002. Extensive cytotoxic lesions of the rat retrosplenial cortex reveal consistent deficits on tasks that tax allocentric spatial memory. Behav. Neurosci. 116 (1), 85–94.
Vann, S.D., Aggleton, J.P., Maguire, E.A., 2009. What does the retrosplenial cortex do? Nat. Rev. Neurosci. 10, 792–802.
Vass, L.K., Epstein, R.A., 2013. Abstract representations of location and facing direction in the human brain. J. Neurosci. 33 (14), 6133–6142.
Vass, L.K., Epstein, R.A., 2016. Common neural representations for visually guided reorientation and spatial imagery. Cereb. Cortex. http://dx.doi.org/10.1093/cercor/bhv343.
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.
Watson, D.M., Hartley, T., Andrews, T.J., 2014. Patterns of response to visual scenes are linked to the low-level properties of the image. Neuroimage 99, 402–410.
Watson, D.M., Young, A.W., Andrews, T.J., 2016b. Spatial properties of objects predict patterns of neural response in the ventral visual pathway. Neuroimage 126, 173–183.
Watson, D.M., Hymers, M., Hartley, T., Andrews, T.J., 2016a. Patterns of neural response in scene-selective regions of the human brain are affected by low-level manipulations of spatial frequency. Neuroimage 124 (Pt A), 107–117.
Willenbockel, V., Sadr, J., Fiset, D., Horne, G., Gosselin, F., Tanaka, J.W., 2010. Controlling low-level image properties: the SHINE toolbox. Behav. Res. Methods 42 (3), 671–684.
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 (4), 679–691.
Wixted, J.T., Squire, L.R., 2011. The medial temporal lobe and the attributes of memory. Trends Cogn. Sci. 15 (5), 210–217.
Wood, E.R., Dudchenko, P.A., Robitsek, R.J., Eichenbaum, H., 2000. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron 27 (3), 623–633.
Wright, P., Randall, B., Clarke, A., Tyler, L.K., 2015. The perirhinal cortex and conceptual processing: effects of feature-based statistics following damage to the anterior temporal lobes. Neuropsychologia 76, 192–207.