

Dynamic Changes In The Medial Temporal Lobe During Incidental Learning Of Object–Location Associations

Anna Manelis^{1,2}, Lynne M. Reder^{1,2} and Stephen José Hanson³

¹Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213, USA, ²The Center for the Neural Basis of Cognition, a joint project of Carnegie Mellon University and the University of Pittsburgh, Pittsburgh, PA 15213, USA and ³Department of Psychology, Rutgers, The State University of New Jersey, Newark, NJ 07102, USA

Address correspondence to Anna Manelis, Department of Psychology, Carnegie Mellon University, Baker Hall 342c, Pittsburgh, PA 15213, USA. Email: anna.manelis@gmail.com.

The role of the medial temporal lobe (MTL) in associative memory encoding has been the focus of many memory experiments. However, there has been surprisingly little investigation of whether the contributions of different MTL subregions (amygdala, hippocampus [HPC], parahippocampal [PHc], perirhinal cortex [PRc], and temporal polar cortex [TPc]) shift across multiple presentations during associative encoding. We examined this issue using event-related functional magnetic resonance imaging and a multivoxel pattern classification analysis. Subjects performed a visual search task, becoming faster with practice to locate objects whose locations were held constant across trials. The classification analysis implicated right HPC and amygdala early in the task when the speed-up from trial to trial was greatest. The same analysis implicated right PRc and TPc late in learning when speed-up was minimal. These results suggest that associative encoding relies on complex patterns of neural activity in MTL that cannot be expressed by simple increases or decreases of blood oxygenation level—dependent signal during learning. Involvement of MTL subregions during encoding of object–location associations depends on the nature of the learning phase. Right HPC and amygdala support active integration of object and location information, while right PRc and TPc are involved when object and spatial representations become unitized into a single representation.

Keywords: associative memory encoding, fMRI, multivoxel pattern classification, medial temporal lobe, object–location associations

Introduction

Modern theories of the medial temporal lobe functioning are in agreement that the hippocampus (HPC) is critical for learning of new associations. The HPC binds together distinct pieces of information to form relational representations that are domain general and flexible in nature (e.g., Norman and O'Reilly 2003; Davachi 2006; Diana et al. 2007; Eichenbaum et al. 2007; Squire et al. 2007; Henke 2010). The role of amygdala in binding is not as well defined. While some studies reported greater amygdala activation for processing of items compared with associations (e.g., Achim and Lepage 2005), others found greater involvement of the amygdala in formation of associations compared with processing of items (e.g., Killgore et al. 2000; Kirwan and Stark 2004).

Several recent studies identified perirhinal cortex (PRc) and parahippocampal cortex (PHc) as the regions important for associative learning and memory. For example, PRc was engaged in processing of intraitem associations (Staresina and Davachi 2006, 2008), novel pairs of pictures (Pihlajamäki et al.

2003), and unitized associations (Haskins et al. 2008). PHc was shown to be involved in processing of spatial and nonspatial associations (e.g., Goh et al. 2004; Aminoff et al. 2007; Peters et al. 2009). One recent study reported the anterior/posterior differentiation in the PHc with anterior regions involved in encoding of location information and posterior regions involved in encoding of both location and object information (Buffalo et al. 2006).

Temporal polar cortex (TPc) is critical for semantic memory (e.g., Patterson et al. 2007). The degeneration of TPc results in the semantic dementia, an inability to name and remember properties of objects (e.g., Martin 2007; Patterson et al. 2007). Recent work suggests that this region may also be critical for rapid acquisition of novel associations (Sharon et al. 2011). The subjects in that study learned associations between objects and their names either in the “fast mapping” experimental paradigm or in the explicit encoding paradigm. The subjects with the damage to HPC successfully recognized associations acquired through the fast mapping but not through explicit encoding. Noteworthy, the subject who had spared HPC but impaired TPc performed comparably to controls on the explicit encoding task but performed worse than the controls on the fast mapping task.

While the studies described above suggest the importance of all MTL subregions for associative memory, there has been little discussion of whether the contributions of these MTL subregions might change across multiple presentations during associative encoding. The present study addresses this question using event-related functional magnetic resonance imaging (fMRI) during a visual search task: Subjects search for a target that is repeatedly presented either in a fixed location or in varied locations.

Most of the studies that examined associative memory encoding used univariate methods of neuroimaging data analysis and, as such, only allow identifying the average level of neural activation in a brain region. This raises a question of whether repetition-related changes of neural activation in a specific MTL subregion reflect qualitatively different patterns of activation on 2 consecutive presentations or smaller modifications within the same pattern (e.g., the same neurons have a lower firing rate). While not many researchers have tried to address this question, those who did have found that the activation patterns in MTL are indeed more complex (Hassabis et al. 2009). In the current study, we explore the multivariate nature of neuroimaging data by using a multivoxel classification approach that detects the patterns of activity across multiple voxels (Haxby et al. 2001; Hanson et al. 2004; Hanson and Halchenko 2008; O'Toole et al. 2007). Classification methods

examine “the statistical relationship between patterns of brain activity and the occurrence of particular experimental conditions” (O’Toole et al. 2007, p. 1736). They are more sensitive to the changes in neural activity than the standard univariate methods (e.g., GLM [General Linear Model]) in that they can pick up differences among experimental conditions even when GLM cannot (Diana et al. 2008). This is especially important for studies of HPC, the region that is often difficult to image due to an inherent low signal-to-noise ratio (Greicius et al. 2003; Zeineh et al. 2003).

Hippocampal encoding is usually rapid (e.g., Nakazawa et al. 2004; Bast 2007). Some types of associative encoding (e.g., object-color associations) that engage PRC may also occur in one trial (e.g., Staresina and Davachi 2008). However, other studies indicate that cortical MTL subregions are often engaged over the course of learning (Aminoff et al. 2007; Yassa and Stark 2008; Voss et al. 2009). For example, Voss et al. (2009) compared repeated words (5 repetitions before scanning and 4 presentations in the scanner) to new words and found robust positive correlation between the magnitude of behavioral priming and repetition-related reduction in left PRC. Aminoff et al. (2007) observed a similar effect in PHc for the stimuli that were practiced before suggesting that PHc can be engaged later in learning. These latter findings are consistent with the view that memories that were initially dependent on HPC may become HPC independent later in learning when memory representations formed in HPC are stored in cortical regions (e.g., Eichenbaum 2004; Duff et al. 2006).

The notion that HPC may be important for the maintenance of object-location associations (Ramsøy et al. 2009) and that learning is rapid in HPC suggest that the changes in the patterns of HPC activity during encoding of object-location associations may occur early in learning (e.g., Nakazawa et al. 2004; Bast 2007) and, according to many neuropsychological studies, may be more pronounced in the right hemisphere (e.g., Smith and Milner 1989; Baxendale et al. 1998; Nunn et al. 1999). Similarly, the data recorded from individual MTL neurons in vivo in human epilepsy surgery patients (Rutishauser et al. 2006) suggest that the beginning of learning may also be supported by the amygdala. The fact that Sharon et al. (2011) repeated their associative stimuli only once may be evidence that acquisition of associations also may involve TPC at the beginning of learning.

While we may expect changes in PHc, a region critical for spatial memory (Bohbot et al. 2000; Ploner et al. 2000), changes in PRC are difficult to predict. PRC is mostly involved in the encoding of in-trial associations and within-domain object-object associations (Pihlajamäki et al. 2003; Staresina and Davachi 2006, 2008). Does this mean that PRC is not involved during encoding of between-domain object-location associations? The most intuitive answer is “yes”; however, one recent study (Buffalo et al. 2006) provides evidence for possible involvement of PRC in formation of between-domain associations. In that study, subjects encountered objects in 1 of 16 locations and were instructed to memorize either objects or locations for a subsequent memory test. Buffalo et al. found that PRC was active during both object encoding and location encoding. One explanation for this result is that even though subjects were instructed to memorize just objects or just locations, they instead encoded object-location associations (e.g., Hasher and Zacks 1979). In addition, some studies have shown that object and location components of an association may become unitized (i.e., to form a single representation) during the course of learning (Musen 1996). If this turns out to be the case with our study, then, according to

earlier work (e.g., Haskins et al. 2008), we may find changes in the patterns of neural activity in PRC appear later in learning.

Materials and Methods

Participants

Ten undergraduates ($M = 20$ years, standard deviation = 2.5, 7 females) from Rutgers University participated in this fMRI study for course credit. All subjects were treated in accordance with Rutgers University and University of Medicine and Dentistry of New Jersey Institutional Review Board (IRB) guidelines. They were unaware of the specific design or hypotheses concerning the experiment prior to their participation but were fully debriefed after they completed the experiment.

Design and Procedure

Subjects were scanned during the visual search task. This task required that subjects locate a target object on a grid as quickly and accurately as possible. Figure 1 illustrates a trial in this task. At the beginning of each 8-s period, subjects were shown a target object for 2 s. An empty 4×4 matrix then appeared for 1 s followed by a matrix that contained 12 objects (including the target). These objects always appeared on the perimeter such that the center 4 cells of the matrix were empty. Subjects had 4 s to locate the target in the display and click on the correct object. They responded by clicking on an MRI-compatible trackball. After 4 s, the display cleared regardless of whether or not the object had been located. When a subject clicked on a wrong object within the allowed 4-s period, the display would not disappear indicating that the subject erred and needed to try again. This strategy discouraged subjects from selecting nontarget objects and forced subject accuracy to be 100%. When subjects failed to make a response or used more than one attempt to find the target, the trial was removed from the behavioral and neuroimaging data analyses. If the response was correct, the display cleared and a blank matrix was shown for the rest of the 8-s period. The length of this rest period (intertrial interval) could vary between 1 and 4.5 s.

There were 12 blocks of trials with 8 search trials in each block. All blocks were separated from each other with 16 s of rest. During the rest periods, a “Wait” sign appeared on the screen. Subjects did not make any manual responses during the rest periods but had to stay alert to start the task as soon as the sign Wait disappeared from the screen. Subjects were presented with constant and variable object-location associations. A “constant” object appeared in only one location assigned to that object. A “variable” object could appear in one of several locations that did not include locations assigned to “constant” objects. Each search display contained all 6 “constant” objects and 6 of the 8 “variable” objects. Distractor stimuli in a given display were targets on other trials and were always presented in valid locations for these




Figure 1. Illustration of a trial sequence.

objects. Thus, if an object “A” appeared in the locations “a,” “b,” and “c” as a target, it could never appear in locations “d” or “e” as a distractor. For each block, 3 “constant” and 5 “variable” objects served as targets with the constraint that no object was repeated as a target within a block.

Order of trials within the block as well as the order of the blocks was randomly determined for each subject. Over the course of the experiment, each “constant” object was the target 6 times and each “variable” object was the target 4–12 times. The 8 objects belonging to the variable condition appeared twice in each location. These objects were seen as targets different numbers of times. Two objects appeared as a target 12 times (20% of variable trials), 3 objects appeared as a target 4 times (6.7% of variable trials), and the other 3 objects served as a target on 10, 8, and 6 trials (16.7%, 13.3%, and 10% of variable trials, respectively). This manipulation allowed us to use all 6 variable locations with equal frequency (10 times each) but still vary the display appearance. By manipulating the locations and the presence of variable objects, we made the constant object–location associations less obvious to the subjects. Subjects were neither instructed that some of the object–location pairings would be repeated nor they were informed that they would be tested later on their memory for object–location associations.

Image Acquisition

The event-related fMRI data were acquired using a Siemens 3-T Allegra head-only MR system. At the beginning of the experiment, a high-resolution structural image (time repetition [TR] = 2000 ms, time echo [TE] = 4.38 ms, slice thickness = 1 mm, field of view [FOV] = 220, number of slices = 176, resolution = $0.8594 \times 0.8594 \times 1$) was acquired using a magnetization prepared rapid gradient echo (MP-RAGE) sequence. Functional data (blood oxygenation level—dependent [BOLD] signal) were collected using a gradient echo, echo-planar sequence (TR = 2000 ms, TE = 30 ms, slice thickness = 4 mm, FOV = 220, number of slices = 32, resolution = $3.4375 \times 3.4375 \times 4.0$). A total of 800 volumes were collected during the search task.

fMRI Data Analysis

The images were preprocessed and analyzed with FSL 4.1.5 (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl software. For each raw BOLD data set, we applied nonlinear noise reduction (smallest univalue segment assimilating nucleus), motion correction (MCFLIRT [Jenkinson et al. 2002]), nonbrain removal using BET (Smith 2002), spatial smoothing using a Gaussian kernel of full-width at half-maximum (FWHM) 9 mm, multiplicative mean intensity normalization of the volume at each time point, and high-pass temporal filtering (Gaussian-weighted least-squares straight line fitting, with $\sigma = 25.0$ s). A hemodynamic response function was modeled using a Gamma function.

The 2-stage registration of the low-resolution BOLD images to standard Montreal Neurological Institute (MNI) template was carried out using FLIRT (FMRIB's Linear Image Registration Tool; Jenkinson and Smith 2001; Jenkinson et al. 2002) using the following parameters: a 9-DOF parameter model, normal search ($\pm 90^\circ$), a correlation ratio cost function and trilinear interpolation. First, BOLD images were registered to the high-resolution structural (MPRAGE) images. Second, the high-resolution images were registered to the MNI152_T1_2mm template. Finally, the 2 resulting transformations were concatenated and applied to the original BOLD image (<http://www.fmrib.ox.ac.uk/fsl/flirt/gui.html>) to transform it to the MNI space.

A GLM analysis with target types (constant object location, variable object location, etc.) and target repetitions (presentation 1, presentation 2, etc.) as explanatory variables was conducted using FEAT (FMRIB Expert Analysis Tool). The length of each event (or trial) for a GLM model was calculated as 2 s (target presentation) + 1 s (interstimulus interval) + search RT [Response Time]. The first-level analysis contrasted 2 consecutive presentations for each subject (presentation 1 vs. presentation 2, presentation 2 vs. presentation 3, etc.). Group means for each contrast of interest were computed using ordinary least square mixed effects. Z-statistic images were thresholded at $P < 0.005$ (voxel-wise, uncorrected). While we conducted the whole-brain GLM analyses, we were only interested in the activations

pertaining to the regions of interest (ROIs). Therefore, a corrected threshold for each cluster (i.e., $P_{corrected}$) was determined by Monte Carlo simulation using the AlphaSim program (<http://afni.nimh.nih.gov/pub/dist/doc/manual/AlphaSim.pdf>) with the functional ROI as a mask and FWHM = 9 mm. We will only report the results pertaining to 5 regions: HPC, amygdala, PHC, PRC, and TPC.

Functional Localization

Functional localization of the ROIs was determined using the Harvard-Oxford cortical and subcortical structural probability atlases (developed by the Harvard Center for Morphometric Analysis and are included with the FSL software) (The atlases were created based on the analysis of T_1 -weighted images of 21 healthy male and 16 healthy female subjects [ages 18–50]. The images were individually segmented by the Harvard Center for Morphometric Analysis using semiautomated tools developed in-house, affine-registered to MNI152 space using FLIRT [FSL], and the transforms were then applied to the individual labels. Finally, these maps were combined across subjects to form population probability maps for each label [<http://www.fmrib.ox.ac.uk/fsl/data/atlas-descriptions.html>]). Probabilistic atlases were successfully used in recent studies of MTL functioning (e.g., Lehn et al. 2009; Robinson et al. 2010) and can be considered as an alternative to the individual ROI drawing procedure. Probabilistic atlases are thought to be beneficial for neuroimaging studies because they allow for consistency across studies by eliminating between-study variability in terms of localization of the geometrically or functionally defined ROIs (Robinson et al. 2010). Using these atlases may be especially beneficial for neuroimaging studies that involve “between-subject” multivoxel pattern classification analyses because these analyses often require that subjects have equal number of voxels in the ROIs under investigation.

In probability atlases, each structure is represented as a standard space image with values from 0:100, according to the cross-population probability of a given voxel being in that structure (<http://www.fmrib.ox.ac.uk/fsl/fslview/atlas.html>). In the present study, we used a 50% probability threshold to included voxels in the specific ROI. This procedure enabled each voxel of interest to be assigned to only one ROI (Voss et al. 2009). Right HPC was comprised of 531 voxels, left HPC—501 voxels, right amygdala—278 voxels, left amygdala—227 voxels, right PHC—139 voxels, left PHC—154 voxels, right PRC—311 voxels, and left PRC—272 voxels.

Classification Analysis

A first-level GLM analysis (at the subject level) was used to compute the contrasts between a specific trial type/presentation and baseline (presentation 1 vs. baseline, presentation 2 vs. baseline, etc.). Pairs of resulting contrast of parameter estimates images with the elements of a pair corresponding to the 2 consecutive presentations were registered to the MNI template (2 mm) and then used as inputs to the multivoxel pattern classification analyses. These analyses examined the patterns of activation that distinguish the 2 contrasts that correspond to the 2 consecutive presentations (e.g., [presentation 1 - baseline] vs. [presentation 2 - baseline], [presentation 2 - baseline] vs. [presentation 3 - baseline]) of object–location pairings in the 10 extended MTL subregions (i.e., left and right HPC, amygdala, PHC, PRC, and TPC).

We ran classifiers on both the constant object–location pairs and the variable object–location pairs in order to assure that any success of classification was due to learning the association of the object with its location as opposed to familiarity with the objects or practice with the task. Given our interest is in the constant–location pairs, we restricted our test of variable location pairs to those classification conditions that were above chance for the constant–location pairs. Please note that a failure to classify 2 consecutive presentations does not imply that a specific region was not activated in the task. Rather it means that the activation patterns for 2 presentations are similar (or that the classifier was not sensitive enough to detect existing differences).

The stability of classification performance was ensured by testing the same pairs of trials using 3 classifiers: a sparse multinomial logistic regression (SMLR; Krishnapuram et al. 2005), a linear support vector machine (SVM; Vapnik 1995) on the features (voxels) selected by SMLR and SVM on all features in a specific ROI. The latter was used to explore

the sensitivity of the classifier to sparse feature selection. All classifiers were implemented using the multivariate pattern analysis in python software <http://www.pympva.org>; Hanke, Halchenko, Sederberg, Hanson, et al. 2009; Hanke, Halchenko, Sederberg, Olivetti, et al. 2009). Monte Carlo simulation with 5000 replications (Schaffera and Kima 2007) using R statistical package (<http://www.r-project.org/>) indicated that classification accuracy should be at or above 80% to be considered significant.

SMLR And Nested Cross-Validation

Feature selection is an important step in a classification procedure (e.g., Haxby et al. 2001; Hanson et al. 2004). The SMLR classifier allows optimizing for the number of features in the data set by adjusting the l_1 parameter (Krishnapuram et al. 2005). Such adjustment (or “tuning”) may result in a bias in error estimation and, consequently, in poor generalizability of results (Varma and Simon 2006). To avoid this bias, we used a 2-level nested cross-validation (CV) method. Nested CV is an unbiased procedure to select the optimization parameters for a classifier (in the case of SMLR—the number of features in the data set). Every time a classifier is optimized, there is the danger that the “best” performance is due to chance and that it is specific to a subset of subjects. Nested CV helps to avoid this by first pulling out one subject successively for a global validation test. In our case, we created 10 data sets involving 10 - 1 = 9 subjects in each training data set, with a different subject removed for each one. These training sets are then split again to optimize the number of selected voxels. We refer to this first level of CV as the “inner loop.” The second level of CV, referred to here as the “outer loop,” was used to compute an estimate of an error. The best classification parameter selected through the nested CV procedure is used in the training on the full training data set and is tested against the global test subjects.

In the inner loop, we used 10 subsets of 9 subjects taken from a total set of 10. The SMLR classifier was trained on 8 subjects and tested on 1. Eight optimization parameters ($l_1 = 0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1, 2$; a smaller l_1 parameter corresponds to a larger number of features left in the data set) were examined for each ROI for each comparison condition. The classification accuracies and a number of features in the model were then averaged across 10 subsets of data separately for each l_1 parameter (see Supplementary material). The l_1 parameter with the highest average classification accuracy and the largest average number of features (voxels) left in the data set was chosen as the best parameter and was used to estimate the classification error in the outer loop. The outer loop SMLR used all 10 subjects. The classifier was trained on 9 subjects and tested on 1 (using a leave-one-out CV approach). It was repeated for each ROI for each comparison condition.

Results

Behavioral

Figure 2 plots the mean RT for each successive trial as a function of whether the trials were constant or variable object-location




Figure 2. Mean correct RT for locating targets that appeared in constant locations (Constant) versus variable locations (Variable). Black dots on the plot show a power function fit ($y = 1557.5 \times x^{-0.3}$) to search RT for “constant” targets.

pairs. Subjects were able to locate the target object on the first trial 96% of the time. Only these correct trials were used for the analyses of search RT and fMRI data. Consistent with previous findings (Musen 1996; Manelis et al. 2011), search RTs became faster with successive repetitions of a target provided that it appeared in fixed spatial location, $F_{5,45} = 9.2, P < 0.001$; however, when a target was repeated in variable locations, the RTs did not differ across repetitions, $P > 0.1$. The decrease in search RT across the 6 presentations of an object in a constant location was best fit by a power function, $y = 1557.5 \times x^{-0.3}$, R -square = 0.97 (Anderson 1982; Logan 1988).

Classification Analysis of Neuroimaging Data

The best optimization parameter for SMLR was selected based on the results of nested CV (for more details, see Tables S1–S5 in Supplementary Data). This parameter defined the sparseness of the SMLR classifier (i.e., the number of features in the data set). Please note that the same features that were selected for SMLR were also used for SVM (i.e., SVM on SMLR features). Another SVM classifier was run on the full data set. Figure 3 displays the training and cross-validation accuracy (ACC) for SMLR and SVM classifiers. CV accuracies in all ROIs for all conditions of interest are also presented in Table S6 in Supplementary Data. The training accuracy for classifiers that were run on different comparison conditions for different ROIs ranged from 67% to 100%. However, more than 80% of all classification conditions had a training accuracy that ranged between 90% and 100%. CV accuracy at or above 80% was achieved in right HPC, right amygdala, right PRC, and right TPC.

Right HPC and amygdala exhibited distinctive activation patterns at the beginning of learning that distinguished the first




Figure 3. Training and CV accuracy for SMLR and SVM classifiers. The darker the square, the higher the classification accuracy. The MTL subregions are shown on the x-axis. The classification conditions are shown on the y-axis. “C” stands for a constant object-location association (e.g., C1–C2 refers to classification of the first presentation of a target in a constant location vs. the second presentation of the target in the same location). “R” stands for right, “L” stands for left. Am; amygdala.

from the second presentation. This result was replicated using all 3 classifiers (SMLR, SVM on SMLR features, and SVM on all features) in HPC, SMLR(ACC) = 90%, number of features = 440, SVM on SMLR(ACC) = 80% and SVM on all features = 80%. In the amygdala, SMLR classified the 2 conditions with 80% accuracy (278 features). The SVM classifiers showed lower classification performance (ACC = 70%), which was below our 80% threshold. Right PRc and TPc revealed distinct activation patterns in the end of learning for the fifth compared with the sixth presentation. This pattern was stable across classifiers in the PRc, SMLR(ACC) = 80%, number of features = 32, SVM on SMLR(ACC) = 80% and SVM on all features = 85% but not in the TPc. While SVM on SMLR features classified the 2 presentations in the TPc with 80% accuracy, SMLR classification accuracy was slightly below the significance level (ACC = 75%, number of features = 51) and SVM accuracy was at chance (ACC = 55%).

CV performance for classification of the intermediate (2 vs. 3; 3 vs. 4, and 4 vs. 5) constant object–location trials was below the 80% threshold. CV accuracy in left HPC, left amygdala, and bilateral PHc at the beginning of learning (presentation 1 vs. presentation 2 and presentation 2 vs. presentation 3) was at 65–70%, which is nominally above chance but below the significance level.

In order to test whether high classification performance is due to learning of object–location associations as opposed to merely becoming more familiar with the objects, we ran classifiers on the variable object–location pairings for those

conditions that were successfully classified for constant pairings. For variable object–location pairings, classification was not above chance for either presentation 1 versus 2 in right HPC (SMLR(ACC) = 60%, SVM on SMLR(ACC) = 55% and SVM on all features = 45%) nor for presentation 1 versus 2 in right amygdala (SMLR(ACC) = 40%, SVM on SMLR(ACC) = 35% and SVM on all features = 40%). Presentations 5 and 6 for variable object–location pairings were also classified at chance in right PRc (SMLR(ACC) = 50%, SVM on SMLR(ACC) = 45%, and SVM on all features = 45%) and in right TPc (SMLR(ACC) = 50%, SVM on SMLR(ACC) = 55%, and SVM on all features = 55%).

Comparison Of Bold Signal Changes Between 2 Consecutive Presentations

We examined the BOLD signal increases and decreases in each of ROIs using 2 methods. A conventional univariate GLM approach identified voxels differentially activated for 2 consecutive presentations. The second approach investigated the BOLD signal changes in the voxels that were diagnostic for a specific classification condition as shown by the SMLR classifier.

Univariate Method

The results of a univariate analysis are illustrated in Figure 4A. Greater activation on the first compared with the second presentation of constant object–location associations was revealed in the cluster of 10 voxels in the intersection of right

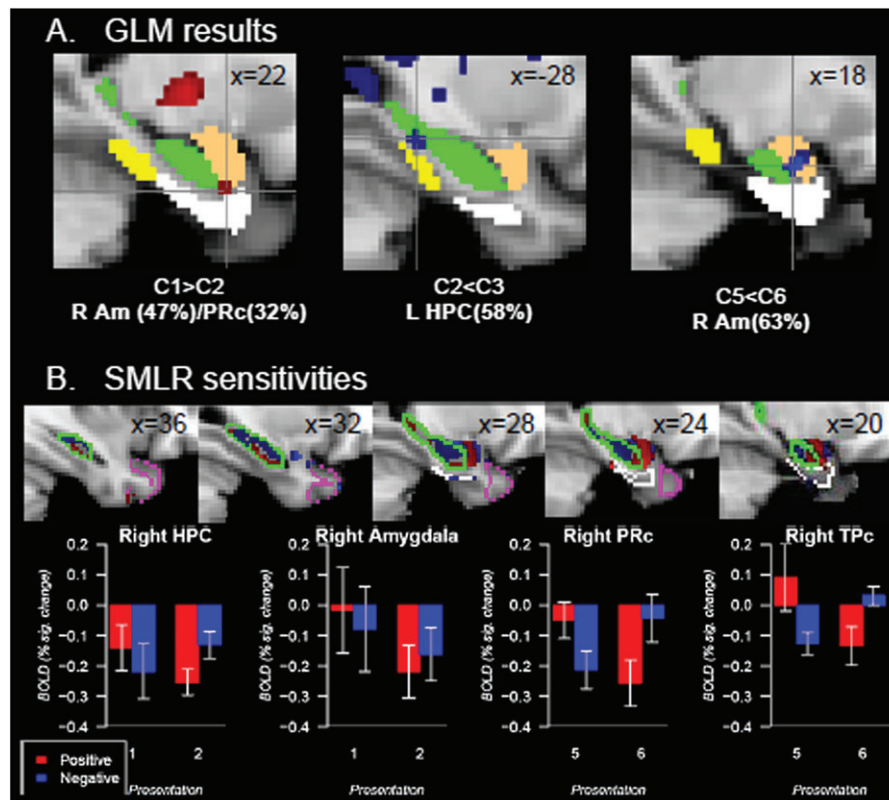


Figure 4. (A) The results of the GLM analysis in right amygdala (R Am) and left hippocampus (L HPC). "C" stands for constant object–location associations. A number next to the "C" indicates the presentation number (e.g., C1 means the first presentation of constant pairings). (B) Top images represent diagnostic sensitivity of voxels to comparison conditions as revealed by the SMLR classifier (C1 vs. C2 in right HPC and amygdala, C5 vs. C6 in right PRc and TPc). A bottom figure illustrates BOLD signal changes relative to a resting baseline in right HPC, amygdala, PRc, and TPc as a function of repetition and voxel sensitivity. A tan mask represents amygdala, a green mask and a contour represent HPC, a white mask and a contour represent PRc, a yellow mask represents PHc, a magenta contour represents TPc. All regions are presented according to the Harvard-Oxford cortical and subcortical structural probability atlases with the probability for a voxel to be in a specific region at 50–100%.

amygdala (47%) (The number in the parenthesis refers to a probability of a peak voxel in the cluster to be located in the ROI according to the Harvard-Oxford cortical and subcortical structural probability atlases.), right PRC (32%) and right HPC (15%) with a peak activation at (22, -2, -28), $z = 2.78$, $P_{\text{corrected}} < 0.05$. (Please note that we computed $P_{\text{corrected}}$ using the right amygdala as a mask in the AlphaSim program because the probability of a peak voxel to be in the right amygdala was the largest and the closest to 50% threshold.) There was less activation on the fifth relative to the sixth presentation of constant pairings in the right amygdala (63%) ([18, -6, -20], $z = 3.66$, number of voxels = 16). Less activation on the second relative to the third presentation of constant associations was observed in the left HPC (58%) ([-28, -32, -12], $z = 2.91$, number of voxels = 16, $P_{\text{corrected}} < 0.05$). Finally, right TPC (69%) showed less activation on Presentation 4 compared with Presentation 5 ([30, 18, -32], $z = 2.86$, number of voxels = 10, $P_{\text{corrected}} > 0.1$). Given that the $P_{\text{corrected}}$ is greater than a conventional $\alpha < 0.05$, which increases the probability of a false detection, we will not discuss this result further in the manuscript.

Bold Signal Changes In The Voxels Diagnostic For A Specific Condition

Classifiers provided us with the diagnostic sensitivity to a particular classification condition in each voxel. These positive and negative sensitivities (positive were pertaining the one classification condition and negative were pertaining to another) computed by the SMLR classifier were used to test whether voxels that were differentially sensitive to each of comparison conditions during classification would show distinct increase/decrease patterns for these conditions (Fig. 4B).

A two-way analysis of variance with repetition and sensitivity directionality as repeated measures was conducted on BOLD signal changes for the first versus second presentations in right HPC and right amygdala and for the fifth versus sixth presentations in right PRC. We found no significant main effects but a significant interaction between stimulus repetition and SMLR sensitivity for a given condition (right HPC: $F_{1,9} = 18.0$, $P < 0.005$, right amygdala: $F_{1,9} = 9.5$, $P < 0.05$, right PRC: $F_{1,9} = 43.2$, $P < 0.001$, right TPC: $F_{1,9} = 19.9$, $P < 0.005$). In right HPC, voxels that were diagnostic for the first presentation of object-location associations decreased neural response to the second, relative to the first presentation. In contrast, voxels that were diagnostic for the second presentation increased their neural response to the second, relative to the first presentation. In right amygdala, voxels diagnostic for the first presentation showed greater repetition-related decreases compared with voxels diagnostic for the second presentation. In right PRC, voxels that were diagnostic for the fifth presentation of object-location associations showed repetition-related decreases while voxels sensitive to the sixth presentation showed repetition-related increases. A lack of significant main effects as well as a low magnitude of changes between comparison conditions in specified voxels suggests why we often fail to detect significant differences using GLM.

Discussion

The goal of this study was to examine whether and how involvement of HPC, amygdala, PHc, PRC, and TPC change over

the course of learning in an incidental, spatial-localization task. If a specific MTL subregion is dynamically involved in associative memory encoding, then it is reasonable that the patterns of neural activity corresponding to the 2 consecutive presentations of associations will differ. By using SMLR and SVM classifiers, we found evidence that involvement of different MTL regions changes during the course of incidental learning of object-location pairs. We will discuss these findings and their implications below for each of the individual MTL subregions.

Hippocampus and Amygdala

Previous work has implicated HPC in binding of new object-location associations (e.g., Burgess et al. 2002). HPC involvement in cognitive tasks was rarely explored using a classification approach; however, one recent study suggests that this method may help to reveal novel information about HPC functioning (Hassabis et al. 2009). Similar to our study, Hassabis et al. (2009) used a multivariate pattern classification approach to decode the patterns of neural activity in MTL. The difference between our 2 studies is that their task involved spatial navigation rather than visual search. Hassabis et al. (2009) found that a subject's positions in a room were accurately predicted by many voxels in the posterior bilateral HPC but by only a few voxels in cortical MTL subregions.

The present study extends these findings and offers new insights into the nature of spatial learning. Specifically, we found that the initial processing of an object in a location relies on "qualitatively" different patterns of activity in right HPC and amygdala than "reencoding" of the same object-location association on the subsequent trial. This result is complemented by the finding from the GLM analysis that revealed decreases in BOLD signal on the second presentation compared with the first presentation in the region located on the border of right amygdala, PRC and HPC, and the finding that the largest behavioral change (i.e., speedup in RT) between trials also occurs between the first and second presentations. Furthermore, the "quantitative" changes were observed in the left HPC with greater activation for Presentation 3 compared with Presentation 2. Taken together, these new findings refine earlier conclusions concerning the involvement of MTL in learning (e.g., Poldrack et al. 2001) suggesting that HPC and the amygdala, but not cortical MTL subregions, are engaged in the early stage of learning. These results also support evidence coming from animal research (e.g., Nakazawa et al. 2004; Bast 2007) and human recordings from individual MTL neurons (Rutishauser et al. 2006) that indicate that learning in HPC and amygdala is rapid and can occur in one trial.

Previous studies implicated the amygdala in regulation of attention (Gallagher and Holland 1994) and showed that the amygdala is more active for trials with a larger memory load (Schon et al. 2009) and during encoding of objects rather than locations (Ramsøy et al. 2009). Our study provides support for the idea that the amygdala is also engaged in associative memory encoding (e.g., Killgore et al. 2000). One recent study (Rutishauser et al. 2006) reported that there are 2 classes of neurons in the amygdala: one class of neurons increases firing rate for novel stimuli, another class of neurons increases firing for familiar stimuli. Our study provides clear support for these findings by showing that the engagement of the amygdala is not limited to the beginning of learning as described above but extends to the later stages of learning as well. Univariate analyses of repeated trials revealed that the right amygdala

significantly increases its activation from Presentation 5 to Presentation 6. In light of Rutishauser et al. (2006) findings, the increase in activation later in learning may indicate that the object–location associations became sufficiently familiar to the subjects by the sixth presentation. It is noteworthy that the classification analyses provided some additional support for this interpretation: SVM on SMLR features was able to distinguish Presentations 5 and 6 in the right amygdala with 70% accuracy which is nominally above chance but below the significance level.

Poor generalization performance does not mean that a region is not involved in learning of object–location associations (especially if a univariate method of analysis shows that it is involved), but it may mean that the patterns of neural activity differ across subjects. If so, then this variation precludes the classifier from finding common patterns of neural activity for a specific classification condition. Arguably, between-subject variability may be one reason for the nonsignificant classification result in the right amygdala late in learning. The same argument may be applied to the left HPC where the GLM detected significant change in the BOLD activation but the classification accuracy was not significantly above chance (65–70%).

Parahippocampal Cortex

Multiple studies have demonstrated the role of PHc in processing of spatial information (e.g., Epstein and Kanwisher 1998; Bohbot et al. 2000; Ploner et al. 2000) and, specifically, object–location associations (e.g., Sommer et al. 2005). Therefore, it was surprising to find that the PHc was not involved in this study. One difference between our study and that of Sommer et al. (2005) is that our subjects learned these pairings incidentally while their subjects were required to intentionally learn object–location associations. Conceivably, the intentional learning of object–location associations requires more PHc processing than incidental binding.

Another explanation for the difference in findings is that our multivariate method involved finding patterns of activation that were not only diagnostic for comparison conditions but also common across all subjects. Therefore, if activation patterns from trial to trial in PHc are not in sync across subjects, there will be poor generalization performance of the classifiers. Some support for this view comes from our finding that, during the early stages of learning, classification accuracy in PHc was 65–70%, which is nominally above chance (50%) but below our 80% threshold.

Perirhinal Cortex

Recent studies suggest that the PRc may be involved in association of in-trait elements of a stimulus (e.g., Staresina and Davachi 2008) and object–object associations (Pihlajamäki et al. 2003). Here, we provide further evidence that the PRc plays a role in associative encoding but suggest that this encoding need not be limited to within-domain associations. While our findings implicated both HPC and PRc in associative memory encoding, consistent with previous studies (e.g., Danckert et al. 2007), the role of PRc was clearly dissociable from that of HPC. In contrast to the right HPC, right PRc became engaged late in learning, when the speedup in time to locate the object had approached asymptote.

In trying to account for these data, there are several views of PRc function that should be considered. A number of studies

have identified PRc involvement in item recognition based on a familiarity process (e.g., Brown and Aggleton 2001; Ranganath et al. 2004). Conceivably then, our PRc data reflect subjects' increased familiarity with the objects they must locate. The problem with this interpretation is that the classification performance was at chance for the variable object–location pairings for the same conditions that were above chance for the constant object–location pairings. Presumably familiarity with the objects, per se, should not depend on the consistency of the search location.

Another view suggests that learning in cortical MTL regions is slow because they integrate experiences over multiple trials (e.g., O'Reilly and Rudy 2001). From this point of view, the change in the pattern of PRc activation later in the task may simply reflect slower learning. The problem with this account is that search RT reached asymptote on the fourth presentation of object–location pairs but the pattern of neural response in PRc changed only on the sixth presentation.

In our view, the qualitative change in the PRc activation pattern late in learning reflects PRc processing of the newly unitized object–location representations. Our view is consistent with a recent study of Staresina and Davachi (2010). They proposed that the process of unitization occurs in non-MTL cortical regions before information reaches the PRc and that PRc activation increases when the representation of a stimulus is enhanced.

The unitization account of PRc functioning fits well with previous findings from patient studies that suggest a general role of cortical regions for encoding of unitized representations. Amnesiacs with damage to HPC were able to learn associations if components of those associations were unitized (e.g., Quamme et al. 2007). In the Duff et al. (2006) study, amnesic patients learned associations between abstract pictures (tangrams) and labels that patients themselves generated for those pictures. These patients were, however, unable to learn associations between tangrams and unrelated words that they had not generated. We suspect that the self-generated labels and the tangrams formed unitized representations that were successfully encoded in the cortical regions rather than in HPC.

The notion that units or chunks are formed with practice has a long history in psychology (e.g., Miller 1956; Chase and Simon 1973; Gobet et al. 2001). As with all other learning processes, formation of unitized representations occur gradually: increasing in strength with practice and losing the strength with disuse, over time. During the course of the experiment, the association of objects to repeated locations gradually strengthened and at some point this ever-stronger representation was sufficiently strong to be processed by the PRc. We believe that searching for objects in constant locations results in their unitization, an idea put forward by Musen (1996).

The behavioral evidence for the strengthening of object–location associations into chunks comes from an additional phase of our experiment, not described in previous sections (but see Manelis et al. 2011). After completing the 12 blocks of visual search in which some objects were always shown in constant locations and other objects were shown in varied locations, the mappings changed such that “constant” objects were now shown in locations that had been reserved for variable location items and “variable” objects were now presented in locations that had been reserved for “constant” objects. Time to search for the former “constant” objects was

significantly slowed when they were presented in new locations ($t_5 = -3.8$, $P < 0.005$). In contrast, the search times for variable location objects were unaffected by the swap, $P > 0.1$. It is likely that a large behavioral cost for location swap in the constant but not variable condition occurred because subjects have unitized constant, but not variable, object–location representations during the first phase of the experiment.

Temporal Polar Cortex

Sharon et al. (2011) reported that patients with damage to the HPC, but intact TPC, were able to acquire novel semantic associations between objects and their names provided they were assigned to the fast mapping condition. In the fast mapping condition, encoding the relationship between features is an automatic aspect of the task and as a consequence novel associations are created within “a pragmatic communication situation and an existing semantic context” (Sharon et al. 2011, p. 1146). From this perspective, our study also may be considered a fast mapping experiment and, therefore, might replicate the Sharon et al. (2011) findings. In other words, if TPC is critical for rapid acquisition of associations, we should observe either quantitative or qualitative changes in this region at the beginning of our task. The fact that neural activation in this region was not modulated early in learning suggests that TPC involvement is specific for acquisition of semantic but not object–location associations. Indeed, in his review of semantic dementia, Patterson et al. (2007) describes Mr. M, a patient with the degeneration in the anterior temporal lobes, who had profound semantic dementia but spared spatial memory that is required for route navigation.

While our study does not support the idea that TPC is critical for formation of new associations, it provides new insights concerning the contributions of this region during multiple presentations of elements that may become associated. Evidence from the study of patients with anterior temporal lobectomy (Glosser et al. 2003) offers an explanation for the modulation in right TPC activation that is observed later in learning. According to Glosser et al. (2003), damage to the right TPC impairs learning of new faces and recognition of familiar ones. Along with the more general finding that both item encoding and retrieval engage right TPC (Persson and Nyberg 2000), this finding suggests the importance of this region for object processing. If, in our study, TPC involvement was related solely to processing of items, we would be able to classify Presentations 5 and 6 not only for constant pairings but also for variable object–location association. As in the case with PRC, we failed to classify variable pairings. Therefore, given the proximity and connectivity of TPC and PRC (Blaizot et al. 2010), it is reasonable to suggest that the change in the activation pattern in the right TPC later in learning could be related to processing of already unitized associations.

Using The Multivariate Pattern Classification Approach As A Method To Study Mtl

It is worth emphasizing that multivariate classification methods and univariate methods of fMRI data analysis (i.e., GLM) often played a complementary role in this study. Sometimes classifiers were able to accurately predict a comparison condition in MTL subregions when univariate methods failed. However, sometimes the classifiers failed to show above chance CV performance when the GLM analysis revealed

repetition-related decreases or increases. These suggest that in order to understand both qualitative and quantitative changes in brain activation, one need to use both multivariate and univariate approaches because the nature of the information provided by the 2 methods may differ. Thus, GLM indicates the magnitude of changes in each voxel as a function of experimental condition. A multivoxel classification analysis reveals different subsets of voxels diagnostic for each condition of interest. These subsets of voxels can be used later for exploration of BOLD signal changes. For example, using results from our classification analysis, we discovered that voxels located in the middle part of the right HPC exhibited the repetition-related decreases, while voxels located in the right posterior and anterior HPC showed an increased activation for repetition of object–location associations.

Conclusions

The take-away message from this study is that learning of object–location associations starts with active integration of object and location information in right HPC and amygdala. Subsequent repetitions strengthen this association on a cortical level, possibly by strengthening the connections between cortical regions. Behaviorally, strengthening is expressed as facilitation of search (decreased RT to locate objects) for the objects that appear in fixed locations. Neurally, this strengthening of object–location associations leads to their unitization as single entities with each representation reflecting the integration of this object–location information. Once the association becomes a unit, it is processed in the right PRC and TPC.

Funding

National Institute of Mental Health at the National Institutes of Health (5T32-MH019983-12); James S. McDonnell Foundation.

Supplementary Material

Supplementary material can be found at: <http://www.cercor.oxfordjournals.org/>

Notes

Conflict of Interest: None declared.

References

- Achim AM, Lepage M. 2005. Neural correlates of memory for items and for associations: an event-related functional magnetic resonance imaging study. *J Cogn Neurosci*. 17:652–667.
- Aminoff E, Gronau N, Bar M. 2007. The parahippocampal cortex mediates spatial and nonspatial associations. *Cereb Cortex*. 17:1493–1503.
- Anderson JR. 1982. Acquisition of cognitive skill. *Psychol Rev*. 89:369–406.
- Bast T. 2007. Toward an integrative perspective on hippocampal function: from the rapid encoding of experience to adaptive behavior. *Rev Neurosci*. 18:253–281.
- Baxendale SA, Thompson PJ, Paesschen WV. 1998. A test of spatial memory and its clinical utility in the pre-surgical investigation of temporal lobe epilepsy patients. *Neuropsychologia*. 36:591–602.
- Blaizot X, Mansilla F, Insausti AM, Constans JM, Salinas-Alamán A, Pró-Sistiaga P, Mohedano-Moriano A, Insausti R. 2010. The human parahippocampal region: I. temporal pole cytoarchitectonic and MRI correlation. *Cereb Cortex*. 20:2198–2212.

- Bhobot VD, Allen JJ, Nadel L. 2000. Memory deficits characterized by patterns of lesions to the hippocampus and parahippocampal cortex. *Ann N Y Acad Sci.* 911:355-368.
- Brown MW, Aggleton JP. 2001. Recognition memory: what are the roles of the perirhinal cortex and hippocampus? *Nat Rev Neurosci.* 2:51-61.
- Buffalo EA, Bellgowan PSF, Martin A. 2006. Distinct roles for medial temporal lobe structures in memory for objects and their locations. *Learn Mem.* 13:638-643.
- Burgess N, Maguire EA, O'Keefe J. 2002. The human hippocampus and spatial and episodic memory. *Neuron.* 35:625-641.
- Chase WG, Simon HA. 1973. Perception in chess. *Cogn Psychol.* 4:55-81.
- Danckert SL, Gati JS, Menon RS, Köhler S. 2007. Perirhinal and hippocampal contributions to visual recognition memory can be distinguished from those of occipito-temporal structures based on conscious awareness of prior occurrence. *Hippocampus.* 17: 1081-1092.
- Davachi L. 2006. Item, context and relational episodic encoding in humans. *Curr Opin Neurobiol.* 16:693-700.
- Diana RA, Yonelinas AP, Ranganath C. 2007. Imaging recollection and familiarity in the medial temporal lobe: a three-component model. *Trends Cogn Sci.* 11:379-386.
- Diana RA, Yonelinas AP, Ranganath C. 2008. High-resolution multi-voxel pattern analysis of category selectivity in the medial temporal lobes. *Hippocampus.* 18:536-541.
- Duff MC, Hengst J, Tranel D, Cohen NJ. 2006. Development of shared information in communication despite hippocampal amnesia. *Nat Neurosci.* 9:140-146.
- Eichenbaum H. 2004. Hippocampus: cognitive processes and neural representations that underlie declarative memory. *Neuron.* 44: 109-120.
- Eichenbaum H, Yonelinas AP, Ranganath C. 2007. The medial temporal lobe and recognition memory. *Annu Rev Neurosci.* 30:123-152.
- Epstein R, Kanwisher N. 1998. A cortical representation of the local visual environment. *Nature.* 392:598-601.
- Gallagher M, Holland PC. 1994. The amygdala complex: multiple roles in associative learning and attention. *Proc Natl Acad Sci U S A.* 91:11771-11776.
- Glosser G, Salvucci AE, Chiaravalloti ND. 2003. Naming and recognizing famous faces in temporal lobe epilepsy. *Neurology.* 61:81-86.
- Gobet F, Lane PC, Croker S, Cheng PCH, Jones G, Oliver I, Pine JM. 2001. Chunking mechanisms in human learning. *Trends Cogn Sci.* 5:236-243.
- Goh JOS, Siong SC, Park D, Gutchess A, Hebrank A, Chee MWL. 2004. Cortical areas involved in object, background, and object-background processing revealed with functional magnetic resonance adaptation. *J Neurosci.* 24:10223-10228.
- Greicius MD, Krasnow B, Boyett-Anderson JM, Eliez S, Schlaggar AF, Reiss AL, Menon V. 2003. Regional analysis of hippocampal activation during memory encoding and retrieval: fMRI study. *Hippocampus.* 13:164-174.
- Hanke M, Halchenko YO, Sederberg PB, Hanson SJ, Haxby JV, Pollmann S. 2009. PyMMPA: a python toolbox for multivariate pattern analysis of fMRI data. *Neuroinformatics.* 7:37-53.
- Hanke M, Halchenko YO, Sederberg PB, Olivetti E, Fründ I, Rieger JW, Herrmann CS, Haxby JV, Hanson SJ, Pollmann S. 2009. PyMMPA: a unifying approach to the analysis of neuroscientific data. *Front Neuroinformatics.* 3:1-13.
- Hanson SJ, Halchenko YO. 2008. Brain reading using full brain support vector machines for object recognition: there is no "face" identification area. *Neural Comput.* 20:486-503.
- Hanson SJ, Matsuka T, Haxby JV. 2004. Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: is there a "face" area? *Neuroimage.* 23:156-166.
- Hasher L, Zacks RT. 1979. Automatic and effortful processes in memory. *J Exp Psychol Gen.* 108:356-388.
- Haskins AL, Yonelinas AP, Quamme JR, Ranganath C. 2008. Perirhinal cortex supports encoding and familiarity-based recognition of novel associations. *Neuron.* 59:554-560.
- Hassabis D, Chu C, Rees G, Weiskopf N, Molyneux PD, Maguire EA. 2009. Decoding neuronal ensembles in the human hippocampus. *Curr Biol.* 19:546-554.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science.* 293:2425-2430.
- Henke K. 2010. A model for memory systems based on processing modes rather than consciousness. *Nat Rev Neurosci.* 11:523-532.
- Jenkinson M, Bannister P, Brady M, Smith S. 2002. Improved optimisation for the robust and accurate linear registration and motion correction of brain images. *Neuroimage.* 17:825-841.
- Jenkinson M, Smith S. 2001. A global optimisation method for robust affine registration of brain images. *Med Image Anal.* 5:143-156.
- Killgore WD, Casasanto DJ, Yurgelun-Todd DA, Maldjian JA, Detre JA. 2000. Functional activation of the left amygdala and hippocampus during associative encoding. *Neuroreport.* 11:2259-2263.
- Kirwan CB, Stark CEL. 2004. Medial temporal lobe activation during encoding and retrieval of novel face-name pairs. *Hippocampus.* 14:919-930.
- Krishnapuram B, Carin L, Figueiredo MAT, Hartemink AJ. 2005. Sparse multinomial logistic regression: fast algorithms and generalization bounds. *IEEE Trans Pattern Anal Mach Intell.* 27:957-968.
- Lehn H, Steffenach H-A, van Strien NM, Veltman DJ, Witter MP, Häberg AK. 2009. A specific role of the human hippocampus in recall of temporal sequences. *J Neurosci.* 29:3475-3484.
- Logan GD. 1988. Toward an instance theory of automatization. *Psychol Rev.* 95:492-527.
- Manelis A, Hanson C, Hanson SJ. 2011. Implicit memory for object locations depends on reactivation of encoding-related brain regions. *Hum Brain Mapp.* 32:32-50.
- Martin A. 2007. The representation of object concepts in the brain. *Annu Rev Psychol.* 58:25-45.
- Miller G. 1956. The magical number seven plus or minus two: some limits on our capacity for processing information. *Psychol Rev.* 63:81-97.
- Musen G. 1996. Effects of task demands on implicit memory for object-location associations. *Can J Exp Psychol.* 50:104-113.
- Nakazawa K, McHugh TJ, Wilson MA, Tonegawa S. 2004. NMDA receptors, place cells and hippocampal spatial memory. *Nat Rev Neurosci.* 5:361-372.
- Norman KA, O'Reilly RC. 2003. Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. *Psychol Rev.* 110:611-646.
- Nunn JA, Graydon FJ, Polkey CE, Morris RG. 1999. Differential spatial memory impairment after right temporal lobectomy demonstrated using temporal titration. *Brain.* 122(Pt 1):47-59.
- O'Reilly RC, Rudy JW. 2001. Conjunctive representations in learning and memory: principles of cortical and hippocampal function. *Psychol Rev.* 108:311-345.
- O'Toole AJ, Jiang F, Abdi H, Pénard N, Dunlop JP, Parent MA. 2007. Theoretical, statistical, and practical perspectives on pattern-based classification approaches to the analysis of functional neuroimaging data. *J Cogn Neurosci.* 19:1735-1752.
- Patterson K, Nestor PJ, Rogers TT. 2007. Where do you know what you know? The representation of semantic knowledge in the human brain. *Nat Rev Neurosci.* 8:976-987.
- Persson J, Nyberg L. 2000. Conjunction analysis of cortical activations common to encoding and retrieval. *Microsc Res Tech.* 51:39-44.
- Peters J, Daum I, Gizewski E, Forsting M, Suchan B. 2009. Associations evoked during memory encoding recruit the context-network. *Hippocampus.* 19:141-151.
- Pihlajamäki M, Tanila H, Hänninen T, Könönen M, Mikkonen M, Jalkanen V, Partanen K, Aronen HJ, Soininen H. 2003. Encoding of novel picture pairs activates the perirhinal cortex: an fMRI study. *Hippocampus.* 13:67-80.
- Ploner CJ, Gaymard BM, Rivaud-Péchéux S, Baulac M, Clémenceau S, Samson S, Pierrot-Deseilligny C. 2000. Lesions affecting the parahippocampal cortex yield spatial memory deficits in humans. *Cereb Cortex.* 10:1211-1216.
- Poldrack RA, Clark J, Paré-Blagoev EJ, Shohamy D, Moyano JC, Myers C, Gluck MA. 2001. Interactive memory systems in the human brain. *Nature.* 414:546-550.
- Quamme JR, Yonelinas AP, Norman KA. 2007. Effect of unitization on associative recognition in amnesia. *Hippocampus.* 17:192-200.

- Ramsøy TZ, Liptrot MG, Skimminge A, Lund TE, Sidaros K, Christensen MS, Baaré W, Paulson OB, Jernigan TL. 2009. Regional activation of the human medial temporal lobe during intentional encoding of objects and positions. *Neuroimage*. 47:1863-1872.
- Ranganath C, Yonelinas AP, Cohen MX, Dy CJ, Tom SM, D'Esposito M. 2004. Dissociable correlates of recollection and familiarity within the medial temporal lobes. *Neuropsychologia*. 42:2-13.
- Robinson JL, Laird AR, Glahn DC, Lovallo WR, Fox PT. 2010. Metaanalytic connectivity modeling: delineating the functional connectivity of the human amygdala. *Hum Brain Mapp*. 31:173-184.
- Rutishauser U, Mamelak AN, Schuman EM. 2006. Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. *Neuron*. 49:805-813.
- Schaffera JR, Kima MJ. 2007. Number of replications required in control chart Monte Carlo simulation studies. *Commun Stat Simul Comput*. 36:1075-1087.
- Schon K, Quiroz YT, Hasselmo ME, Stern CE. 2009. Greater working memory load results in greater medial temporal activity at retrieval. *Cereb Cortex*. 19:2561-2571.
- Sharon T, Moscovitch M, Gilboa A. 2011. Rapid neocortical acquisition of long-term arbitrary associations independent of the hippocampus. *Proc Natl Acad Sci U S A*. 108:1146-1151.
- Smith ML, Milner B. 1989. Right hippocampal impairment in the recall of spatial location: encoding deficit or rapid forgetting? *Neuropsychologia*. 27:71-81.
- Smith S. 2002. Fast robust automated brain extraction. *Hum Brain Mapp*. 17:143-155.
- Sommer T, Rose M, Weiller C, Büchel C. 2005. Contributions of occipital, parietal and parahippocampal cortex to encoding of object-location associations. *Neuropsychologia*. 43:732-743.
- Squire LR, Zola-Morgan J, Clark RE. 2007. Recognition memory and the medial temporal lobe: a new perspective. *Nat Rev Neurosci*. 8:872-883.
- Staresina BP, Davachi L. 2006. Differential encoding mechanisms for subsequent associative recognition and free recall. *J Neurosci*. 26:9162-9172.
- Staresina BP, Davachi L. 2008. Selective and shared contributions of the hippocampus and perirhinal cortex to episodic item and associative encoding. *J Cogn Neurosci*. 20:1478-1489.
- Staresina BP, Davachi L. 2010. Object unitization and associative memory formation are supported by distinct brain regions. *J Neurosci*. 30:9890-9897.
- Vapnik V. 1995. *The nature of statistical learning theory*. New York: Springer.
- Varma S, Simon R. 2006. Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics*. 7:91.
- Voss JL, Hauner KKY, Paller KA. 2009. Establishing a relationship between activity reduction in human perirhinal cortex and priming. *Hippocampus*. 19:773-778.
- Yassa MA, Stark CEL. 2008. Multiple signals of recognition memory in the medial temporal lobe. *Hippocampus*. 18:945-954.
- Zeineh MM, Engel SA, Thompson PM, Bookheimer SY. 2003. Dynamics of the hippocampus during encoding and retrieval of face-name pairs. *Science*. 299:577-580.