Cerebral Cortex
doi:10.1093/cercor/bhr151

Anna Manelis[1,2], Lynne M. Reder[1,2] and Stephen José Hanson[3]

[1]Department of Psychology, Carnegie Mellon University, Pittsburgh, PA 15213, USA, [2]The Center for the Neural Basis of Cognition, a joint project of Carnegie Mellon University and the University of Pittsburgh, Pittsburgh, PA 15213, USA and [3]Department 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.

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

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

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)

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 × 0.8594 × 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 × 3.4375 × 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-*D* *F* parameter model, normal search (±90°), 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 (FMRI 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. -statistic images were thresholded at < 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., [1esbyhtd(in2(gCench)-4n2(g(means4s)-Morphome430(p(cai

edinT6tepr11encthe43[(-tude,the22rew(inT9-4duare)35b.oae)

the sensitivity of the classifier to sparse feature selection. All classifiers were implemented using the multivariate pattern analysis in python software http://www.pymvpa.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.

ML A N C -

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 lm 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 (lm = 0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1, 2; a smaller lm 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 lm parameter (see Supplementary material). The lm 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.

Behavioral

Figure 2 plots the mean RT for each successive trial as a function of whether the trials were constant or variable object–location 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, $F_{5,45} = 9.2$, < 0.001; however, when a target was repeated in variable locations, the RTs did not differ across repetitions, > 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, $= 1557.5 \times ^{-0.3}$, -square = 0.97 (Anderson 1982; Logan 1988).

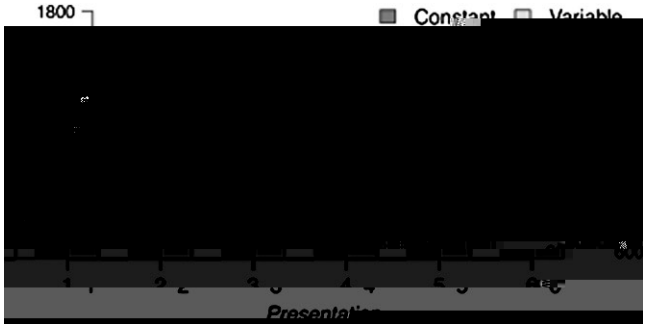

. Mean correct RT for locating targets that appeared in constant locations (Constant) versus variable locations (Variable). Black dots on the plot show a power function fit ($= 1557.5 \times ^{-0.3}$) to search RT for "constant" targets.

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

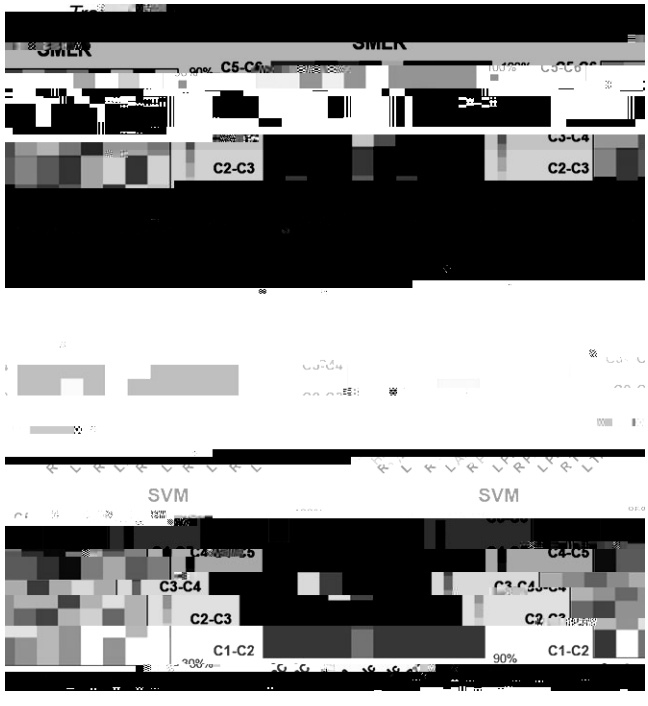

. Training and CV accuracy for SMLR and SVM classifiers. The darker the

from the second presentation. This result was replicated using

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), = 2.78, $_{\text{corrected}}$ < 0.05. (Please note that we computed $_{\text{corrected}}$ 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], = 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], = 2.91, number of voxels = 16, $_{\text{corrected}}$ < 0.05). Finally, right TPc (69%) showed less activation on Presentation 4 compared with Presentation 5 ([30, 18, –32], = 2.86, number of voxels = 10, $_{\text{corrected}}$ > 0.1). Given that the $_{\text{corrected}}$ is greater than a conventional α < 0.05, which increases the probability of a false detection, we will not discuss this result further in the manuscript.

B C I D F A
C

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. 4*B*).

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, < 0.005, right amygdala: $F_{1,9}$ = 9.5, < 0.05, right PRc: $F_{1,9}$ = 43.2, < 0.001, right TPc: $F_{1,9}$ = 19.9, < 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 380(pe)6(r)-5(-655(

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 that 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 intraitem 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

significantly slowed when they were presented in new

Bohbot 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.