High-Resolution Multi-Voxel Pattern Analysis of Category Selectivity in the Medial Temporal Lobes

Rachel A. Diana[1,2,*], **Andrew P. Yonelinas**[2], and **Charan Ranganath**[1,2]

1*Center for Neuroscience, University of California, Davis, CA*

2*Department of Psychology, University of California, Davis, CA*

Abstract

Although the parahippocampal cortex (PHc) is known to be critical for memory formation, little is known about what is encoded by this area. Using multi-voxel pattern analysis of high-resolution functional magnetic resonance imaging (MRI) data, we examined responses to blocks of categorically coherent stimuli and found that patterns of activity in PHc were selective for not only scenes, but also for other nonspatial object categories (e.g., faces and toys). This pattern of results was also found in the parahippocampal place area (PPA), indicating that this region is not sensitive exclusively to scenes. In contrast, neither the hippocampus nor perirhinal cortex (PRc) were found to be selective for category information. The results indicate that regions within the medial temporal lobe may support distinct functions, and that the PHc appears to be particularly sensitive to category-level information.

INTRODUCTION

Medial temporal lobe (MTL) structures are critical for the formation of declarative memories (Eichenbaum et al., 1992), and recent research has focused on specifying the roles of these subregions in the encoding and retrieval of declarative memories. Although the role of the hippocampus has been the focus of much recent research (e.g., Giovanello et al., 2003), the functional role of the cortex surrounding the hippocampus remains relatively unexplored. There is evidence that one of these areas, the parahippocampal cortex (PHc) may be particularly important for encoding information that is spatial in nature (Bohbot et al., 1998; Janzen and van Turennout, 2004). In fact, some data suggest that a specific region within PHC, the parahippocampal place area (PPA), is specialized for encoding of scenes as opposed to objects from other categories (e.g., faces) (Epstein and Kanwisher, 1998; Epstein et al., 1999). However, PHc activity has also been associated with the retrieval of nonspatial source information such as the color of a previously presented item (Ranganath et al., 2003; Diana et al., 2007), as well as with general contextual information associated with an object or event (Bar and Aminoff, 2003). In addition, PHc was found to be more active during viewing of famous faces than unfamiliar faces, providing evidence that the nonspatial contextual information associated with famous faces may be encoded in PHc (Bar et al., in press).

*Correspondence to: Rachel Diana, Center for Neuroscience, University of California, Davis 1544 Newton Court Davis, CA 95616, USA. E-mail: radiana@ucdavis.edu.

To understand the functional characteristics of PHc, it is necessary to understand the nature of neural selectivity in the region. As an initial step toward this goal, the current experiment used high-resolution (voxel size = $1.96 \times 1.96 \times 2$ mm) functional MRI (fMRI) along with a newly developed analysis technique-multi-voxel pattern analysis (MVPA)—to examine the nature of neural selectivity within subregions of the MTL, including PHc, perirhinal cortex (PRc), and the hippocampus. MVPA assesses the pattern of activity across voxels in a given region and uses this pattern to train a neural network model to recognize trials of different types (Haxby et al., 2001; Polyn et al., 2005; Norman et al., 2006). The model is then tested in order to determine whether it generalizes to a new dataset (Norman et al., 2006). Accurate classification in the generalization test suggests that a population of neurons within the region can distinguish between the categories of time points that were analyzed (see detailed methods for further description of MVPA techniques). In the current experiment, eight participants were scanned while viewing blocks of images of different types: scenes, objects, abstract 3D shapes, faces, and toys (see Fig. 1A). The object category was further divided between objects with strong contextual associations (e.g., fire hydrant, traffic light) and objects with weak contextual associations (e.g., banana, shirt) (Bar and Aminoff, 2003). The stimuli were presented to eight participants as blocks of 9 images from a single category with a 10 s delay between each block. Two blocks for each category were presented in each run. Participants performed a "1-back" task, indicating whether each stimulus was identical to the previous stimulus (one repeated image occurred per block). Participants also viewed blocks of scrambled images from the same categories (with spatial frequency phase information scrambled and all other information held constant) in order to assess the type of visual information that was critical for successful classification.

Separate bilateral regions of interest (ROIs) were defined for the hippocampus, PHc, and PRc (see detailed methods for ROI definition guidelines, mean number of voxels per ROI, and mean signal-to-noise (S/N) ratios [SNR] for each ROI). Figure 1B shows sample ROIs for a single subject. Initial analyses using standard univariate general linear modeling (GLM) techniques (Worsley and Friston, 1995) replicated previous reports demonstrating that PHc activation is greater for scenes than for objects, $t(7) = 3.13$, $P < 0.05$ (Epstein and Kanwisher, 1998), and greater for objects with strong contextual associations than for objects with weak associations $t(7) = 3.92$, $P < 0.01$ (Bar and Aminoff, 2003). In the PHc ROI, all pairwise contrasts that involved scenes were significant (face $P < 0.05$; toy $P < 0.01$; abstract shape $P < 0.01$), but no other contrasts were significant (all $P > 0.5$). For the hippocampus and PRc ROIs, none of the contrasts revealed significant effects (all $P > 0.05$).

The MVPA results for objects vs. scenes classification were consistent with the GLM findings described earlier. Accuracy for this classification in PHc was above chance for all participants (see Fig. 1C). Across the group, mean classification accuracy was 69%, a level significantly greater than chance, $t(7) = 7.24$, $P < 0.001$. In contrast to PHc, classification accuracy in the hippocampus was not greater than chance ($M = 50\%$), and was only slightly better than chance in PRc ($M = 55\%$), $t(7) = 2.83$, $P < 0.05$. The classification accuracy measure described above is based on the idea that when the correct output unit in the trained model is more active than the incorrect output unit-even if only by a small amount-the response is considered correct. However, the degree to which the correct output unit is more active than the incorrect output unit may vary from 0.99 to 0.01. Thus, we also characterized the robustness of each network model using a decision confidence measure that assesses how strongly the network settled on the correct response. This measure is based on the activation of the output units (e.g., object and scene for this classification). Table 1 reports the degree to which the correct output unit was more active than the incorrect output unit for each classification. Although PRc produced correct classification of objects vs. scenes (see Table 1A), with the object output unit being more active than the scene output unit for object trials and vice versa, decision confidence was low. We also compared weak vs. strong context stimuli using the MVPA technique. The

analysis did not indicate predictive power in PHc voxels for differentiating weak context object viewing from strong context object viewing ($M = 49\%$). Classification was also at chance for the hippocampus ($M = 53\%$) and PRc ($M = 52\%$) for this comparison. Thus the MVPA results replicate the GLM results for the object vs. scene analysis, but do not replicate the weak context object vs. strong context object results.

We assessed whether PHc is selective for other types of information in addition to scenes by determining whether PHc voxels could differentiate between scenes and objects from other specific categories (faces, toys, and abstract shapes). As shown in Figure 1C, classification accuracy on this four-way discrimination was above chance for every participant, and the group average was significantly greater than chance ($M = 40\%$, where chance is 25%), $t(7) = 10.37$, $P < 0.001$. Classification confidence in PHc (see Table 1B) was greater for the scene output unit than for the specific object output units, although all of the output activations reflected accurate classification. In contrast, accuracy for this classification did not exceed chance levels for the hippocampal ($M = 26\%$) or PRc ($M = 27\%$) ROIs. A further analysis was conducted to determine whether the inclusion of scenes in the four-way classification was entirely responsible for performance being greater than chance. This was not the case, as voxels in the PHc mask predicted face vs. toy vs. abstract shape viewing at a level significantly greater than chance ($M = 42\%$, where chance is 33%), $t(7) = 6.07$, $P < 0.001$. Again, all participants showed classification accuracy that was greater than chance for this comparison. Examination of the confusion matrix (which indicates which stimulus categories were correctly classified) revealed that all three types of stimuli were successfully classified (face $M = 42\%$, toy $M = 46\%$, abstract shape $M = 55\%$). Thus, the three-way object classification did not rely solely on any one object category being successfully classified. Analyses of decision confidence measures also indicated that the individual object categories were successfully classified in PHc and PPA (see Table 1B and 1C). This three-way classification could not be reliably performed with data from the hippocampal ($M = 34\%$) or PRc (37%) ROIs, as accuracy for these regions did not differ from chance levels.

The finding that PHc can successfully classify individual object categories suggests that PHc is not exclusively neurally selective for scenes. However, it is possible that the PPA, a smaller area within PHc, might be neurally selective specifically for scenes. To test this possibility, classification analyses were also conducted on a PPA mask, defined functionally using the same criteria as Epstein and Kanwisher (1998) and contained within the larger PHc ROI. The classification analyses described earlier were repeated for the PPA ROI, revealing above-chance accuracy for objects vs. scenes ($M = 72\%$, chance = 50%), $t(7) = 12.03$, $P < 0.001$, scenes vs. faces vs. toys vs. abstract shapes ($M = 40\%$, chance = 25%), $t(7) = 9.59$, $P < 0.001$, and faces vs. toys vs. abstract shapes ($M = 41\%$, chance = 33%), $t(7) = 4.07$, $P < 0.01$. Thus, the pattern of results for the PPA was identical to that of the PHc, although overall classification accuracy was slightly higher for the PPA ROI. Therefore, we conclude that PPA is not exclusively sensitive to scenes, but is also sensitive to nonspatial information. In fact, comparison of the classification confidence for PHc and PPA (see Table 1A) indicates that PPA produced more confident classification of general objects when compared with scenes. However, in the classifications that included multiple object categories, PHc produced more confident classification of the specific object categories.

Pattern classification is a powerful technique that can potentially use any information available during a trial to distinguish between different categories of stimuli (Cox and Savoy, 2003). To narrow down the type of information that was critical for success in the classifications reported here, we attempted to classify the phase-scrambled images. Spatial frequency phase information is necessary and sufficient for object recognition. If images with scrambled phase information were accurately classified, this would suggest that spatial frequency phase information was not necessary for successful classification. Therefore successful classification

of stimuli with scrambled phase information would indicate that other types of visual information were sufficient for classification. Thus, we repeated the classification analyses described above with scrambled versions of the stimuli from each category. None of these analyses revealed above-chance classification in any of the ROIs, except for comparisons that involved the abstract shapes in PHc and PPA. For example, classification of faces vs. toys vs. abstract shapes in PHc was greater than chance ($M = 39\%$), $t(7) = 3.84$, $P < 0.01$. Closer examination of the confusion matrices (identifying which stimuli produce errors) revealed that only the scrambled abstract shape stimuli were successfully classified ($M = 39\%$), indicating that characteristics of these scrambled stimuli could be distinguished from the other two object categories. However, the scrambled face ($M = 33\%$) and scrambled toy ($M = 32\%$) stimuli could not be classified. It is unclear why scrambled abstract shape stimuli could be accurately classified, but this may be due to the fact that these stimuli consisted of only two colors (black and the color of the shape). Thus the phase-scrambled abstract shapes might be easily distinguishable from other categories, which include a range of colors. Although the abstract shape stimuli can be classified in PHc and PPA based on visual information other than spatial frequency phase information, it seems that phase information is necessary to classify the other categories of stimuli in PHc and PPA.

Overall, the results suggest that the MVPA technique can be more sensitive than standard univariate GLM approaches, as we found that patterns of activity in the PHc ROI could differentiate between categories of nonscene stimuli even though overall activity in the ROI did not significantly differ. However the MVPA technique also provided qualitatively different information than the GLM analysis. For example, weak and strong context objects were not distinguishable using MVPA, even though GLM analyses indicated a significant difference in overall activation between these categories in PHc. This may be due to the type of information that is used in the two analysis techniques. The GLM analysis was based on the average activation of all the voxels in the ROI whereas the MVPA analysis is based on the pattern of voxel activations in the ROI. That is, the pattern classifier relied on particular voxels being activated or deactivated to identify a pattern while the GLM analyses relied on the mean level of activation in the ROI to be greater in one condition than another (regardless of the pattern). If the pattern of voxel activation varied substantially during viewing of successive objects in a "strong context" block, then this might result in poor classification, even if the average activation level averaged across voxels and trials was higher for strong-than for weak-context objects.

The finding that the PPA classifies more accurately than any other ROI is somewhat surprising, given that the PPA mask includes a much smaller region of cortex than the PHc mask, and that the region has the lowest SNR and the smallest number of voxels of any region examined (see detailed methods). Given the robust classification accuracy in the PPA, it is unlikely that poor classification accuracy in the hippocampus and PRc could be attributed to low SNR or smaller numbers of voxels. Instead, it appears that neural coding in the PHc and the PPA is sensitive to differences between the categories of stimuli presented here, whereas coding in the PRc and hippocampus is relatively insensitive to these differences.

Although we did not find category selectivity in the hippocampus, previous research has suggested that individual neurons in the human hippocampus are category selective (Kreiman et al., 2000). However, this single-unit recording study reported category selectivity in only 17 out of 91 neurons tested in the hippocampus. Thus, the population of category-selective neurons in the hippocampus may be too sparse to be accurately measured even using high-resolution fMRI. An alternate possibility is that representations in the hippocampus (Quiroga et al., 2005) and PRc (Murray and Bussey, 1999) may differentiate among individual stimuli. This possibility could be tested in a future study by using MVPA on data from the hippocampus

and PRc during repeated viewing of a particular exemplar from each category, as opposed to multiple exemplars from the same category.

Although the current results do not invalidate the proposal that the PHc is involved in encoding spatial layout information, it is clear that PHc and PPA are not purely selective for spatial information. The spatial hypothesis cannot account for the fact that PHc can differentiate between categories of visual stimuli that lack information about spatial layouts. An alternative possibility is that PHc may encode the general context of each trial (both spatial and nonspatial information). The BIC model of memory (Diana et al., 2007) claims that PHc is involved in processing of context information, while PRc processes item information and the hippocampus binds item and context into an episode. Viewing a series of objects from the same category, organized as a block, may produce a semantic gist that is consistent across the trials of that block. Thus PHc selectivity for "categorical coherence" may be based on the encoding of category-specific details or on the mental set that is associated with encoding multiple items from the same category. Although this study does not confirm or disconfirm the context hypothesis for PHc, future research can test whether category-level information contributes to the context of the image in this type of experiment.

The present results suggest that MVPA, in concert with high-resolution FMRI techniques, can be used to identify reliable patterns of neural coding in the MTL. Future development of this technique may allow researchers to use patterns of MTL activity to track recollection of specific pieces of information in real time (Polyn et al., 2005).

DETAILED METHODS

Strong and weak context stimuli were provided by Moshe Bar and Elissa Aminoff (Bar and Aminoff, 2003). General object stimuli (as well as toys) were gathered from the Amsterdam library of object images (Geusebroek et al., 2005) (available at http://staff.science.uva.nl/~aloi/). Face stimuli were gathered from the AR face database (Martinez and Benavente, 1998) (available at http://cobweb.ecn.purdue.edu/~aleix/aleix_face_DB.html). Abstract shape stimuli were generated specifically for this experiment using the POV-RAY program (available at http://www.povray.org/). Scrambled stimuli were created using Matlab. Scene stimuli were royalty-free images with transferable copyrights gathered from various websites.

MRI data were collected on a 3T Siemens scanner at the UC Davis Research Imaging Center. Functional imaging was done with a gradient echo echoplanar imaging (EPI) sequence optimized for high-resolution functional imaging (TR = 2000 ms, TE = 27, flip angle = 90, FOV = 220 mm, matrix size = 112×112, voxel size = $1.96 \times 1.96 \times 2$ mm) with each volume consisting of 27 contiguous slices. The slices were aligned parallel to the axis of the hippocampus such that the center slice passed through the center of the hippocampus. Additionally, T1-weighted images coplanar with the EPIs were acquired using an MP-RAGE sequence (matrix size = 256×256, voxel size = $1 \times 1 \times 1$ mm, number of slices = 192). fMRI data processing for all subjects included: motion correction using a six-parameter, rigid-body transformation algorithm provided by Statistical Parametric Mapping (SPM5) software (Wellcome Department of Cognitive Neurology) and calculation of the global signal and power spectrum for each scanning run. Data were not spatially smoothed.

Bilateral ROIs corresponding to the hippocampus, PHc, and PRc were defined for each subject using criteria described in previous studies (Insausti et al., 1998; Buffalo et al., 2006). The PHc and PRc were defined as the cortex on both banks of the collateral sulcus. The rostral limit of the PRc was defined 2 mm anterior to the limen insula. The caudal limit of the PHc was defined as the caudal limit of the hippocampus. The border between the PRc and PHc was defined

according to the caudal limit of the head of the hippocampus (the uncus). The hippocampal mask included the hippocampal head, body, and tail (Pruessner et al., 2000). The rostral limit of the hippocampus was defined as the border with the amygdala and the caudal limit was defined as the border with white matter and the tail of the lateral ventricle. The coronal image was used as the default view for labeling of all three regions with reference to sagittal and axial images as needed. The bilateral ROIs for each subregion were also divided into separate left and right hemisphere ROIs in order to determine whether there were any qualitative differences between the two hemispheres. Table 2 reports the mean number of voxels in the original masks as well as the mean number of voxels included for classification. S/N ratio reflects the ratio of the average signal intensity in each ROI over the standard deviation (SD) of the noise in the entire image (Maubon et al., 1999).

BOLD responses were analyzed using a modified general linear model (Worsley and Friston, 1995) as implemented in the VoxBo software package (available at http://www.voxbo.org). Separate covariates modeled BOLD responses to each type of block (e.g., objects, scenes, faces, etc.) as boxcar functions, each convolved with a canonical hemodynamic response function (based on an average of empirically derived hemodynamic response functions from 22 subjects). Each analysis incorporated empirical estimates of intrinsic temporal autocorrelation (see Zarahn et al., 1997 analyses of unsmoothed data) and filters to attenuate frequencies above 0.25 Hz and below 0.001 Hz.

The MVPA toolbox for Matlab, created by the Princeton computational memory lab (available at http://www.csbmb.princeton.edu/mvpa/) was used to conduct pattern classification analyses (Norman et al., 2006). Each classification consisted of the following steps: (1) loading the appropriate ROI mask, (2) loading the EPI data by filtering out voxels that were not selected in the mask, (3) loading a regressor matrix that specifies the task condition corresponding to each image, (4) loading a "runs selector" vector to specify the run in which each TR occurred, (5) z-transforming the data within each voxel relative to its mean value across the time series, (6) creating cross-validation indices to allow classification training and testing based on all possible block combinations (rest TRs were excluded at this point), (7) conducting an analysis of variance (ANOVA) on the data to remove uninformative voxels (thresholded at $P = 0.05$ and run only on training voxels, not testing voxels), and (8) creating a neural network with no hidden layer. The neural network was then trained on a subset of the data (all blocks minus 1) using the "traincgb" function in the Matlab neural network toolbox (a conjugate gradient algorithm using the Powell-Beale reset method) and tested on the remaining data (the untrained block). A test with the "trainscg" function (a scaled conjugate gradient method) produced similar results. Within each classification, all combinations of training and testing runs were executed with the final result being a mean of these iterations. Each classification was then conducted 10 times, with the results reported here being the mean of those 10 classifications.

References

Bar M, Aminoff E. Cortical analysis of visual context. Neuron 2003;38:347–358. [PubMed: 12718867]

Bar M, Aminoff E, Ishai A. Famous faces activate contextual associations in the parahippocampal cortex. Cereb Cortex. in press

Bohbot VD, Kalina M, Stepankova K, Spackova N, Petrides M, Nadel L. Spatial memory deficits in patients with lesions to the right hippocampus and to the right parahippocampal cortex. Neuropsychologia 1998;36:1217–1238. [PubMed: 9842767]

Buffalo EA, Bellgowan PSF, Martin A. Distinct roles for medial temporal lobe structures in memory for objects and their locations. Learn Mem 2006;13:638–643. [PubMed: 16980544]

Cox DD, Savoy RL. Functional magnetic resonance imaging (fMRI) "brain reading": Detecting and classifying distributed patterns of fMRI activity in human visual cortex. Neuroimage 2003;19:261–270. [PubMed: 12814577]

Diana RA, Yonelinas AP, Ranganath C. Imaging recollection and familiarity in the medial temporal lobe: A three-component model. Trends Cogn Sci 2007;11:379–386. [PubMed: 17707683]

Eichenbaum H, Otto T, Cohen NJ. The hippocampus–what does it do? Behav Neural Biol 1992;57:2–36. [PubMed: 1567331]

Epstein R, Kanwisher N. A cortical representation of the local visual environment. Nature 1998;392:598–601. [PubMed: 9560155]

Epstein R, Harris A, Stanley D, Kanwisher N. The parahippocampal place area: Recognition, navigation, or encoding? Neuron 1999;23:115–125. [PubMed: 10402198]

Geusebroek JM, Burghouts GJ, Smeulders AWM. The Amsterdam library of object images. Int J Comput Vision 2005;61:101–112.

Giovanello KS, Verfaellie M, Keane MM. Disproportionate deficit in associative recognition relative to item recognition in global amnesia. Cogn Affect Behav Neurosci 2003;3:186–194. [PubMed: 14672155]

Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science 2001;293:2425–2430. [PubMed: 11577229]

Insausti R, Juottonen K, Soininen H, Insausti AM, Partanen K, Vainio P, Laakso MP, Pitkanen A. MR volumetric analysis of the human entorhinal, perirhinal, and temporopolar cortices. Am J Neuroradiol 1998;19:659–671. [PubMed: 9576651]

Janzen G, van Turennout M. Selective neural representation of objects relevant for navigation. Nat Neurosci 2004;7:673–677. [PubMed: 15146191]

Kreiman G, Koch C, Fried I. Imagery neurons in the human brain. Nature 2000;408:357–361. [PubMed: 11099042]

Martinez AM, Benavente R. The AR face database. CVC Technical Report 24. 1998

Maubon AJ, Ferru J, Berger V, Soulage MC, DeGraef M, Aubas P, Coupeau P, Dumont E, Rouanet JP. Effect of field strength on MR images: Comparison of the same subject at 0.5, 1.0, and 1.5 t. Radiographics 1999;19:1057–1067. [PubMed: 10464808]

Murray EA, Bussey TJ. Perceptual-mnemonic functions of the perirhinal cortex. Trends Cogn Sci 1999;3:142–151. [PubMed: 10322468]

Norman KA, Polyn SM, Detre GJ, Haxby JV. Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. Trends Cogn Sci 2006;10:424–430. [PubMed: 16899397]

Polyn SM, Natu VS, Cohen JD, Norman KA. Category-specific cortical activity precedes retrieval during memory search. Science 2005;310:1963–1966. [PubMed: 16373577]

Pruessner JC, Li LM, Series W, Pruessner M, Collins DL, Kabani N, Lupien S, Evans AC. Volumetry of hippocampus and amygdala with high-resolution MRI, three-dimensional analysis software: Minimizing the diescrepancies between laboratories. Cereb Cortex 2000;10:433–442. [PubMed: 10769253]

Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I. Invariant visual representation by single neurons in the human brain. Nature 2005;435:1102–1107. [PubMed: 15973409]

Ranganath C, Yonelinas AP, Cohen MX, Dy CJ, Tom SM, D'Esposito M. Dissociable correlates of recollection and familiarity within the medial temporal lobes. Neuropsychologia 2003;42:2–13. [PubMed: 14615072]

Worsley KJ, Friston KJ. Analysis of fMRI time-series revisited-again. Neuroimage 1995;2:173–181. [PubMed: 9343600]

Zarahn E, Aguirre GK, D'Esposito M. A trial-based experimental design for fMRI. Neuroimage 1997;6:122–138. [PubMed: 9299386]
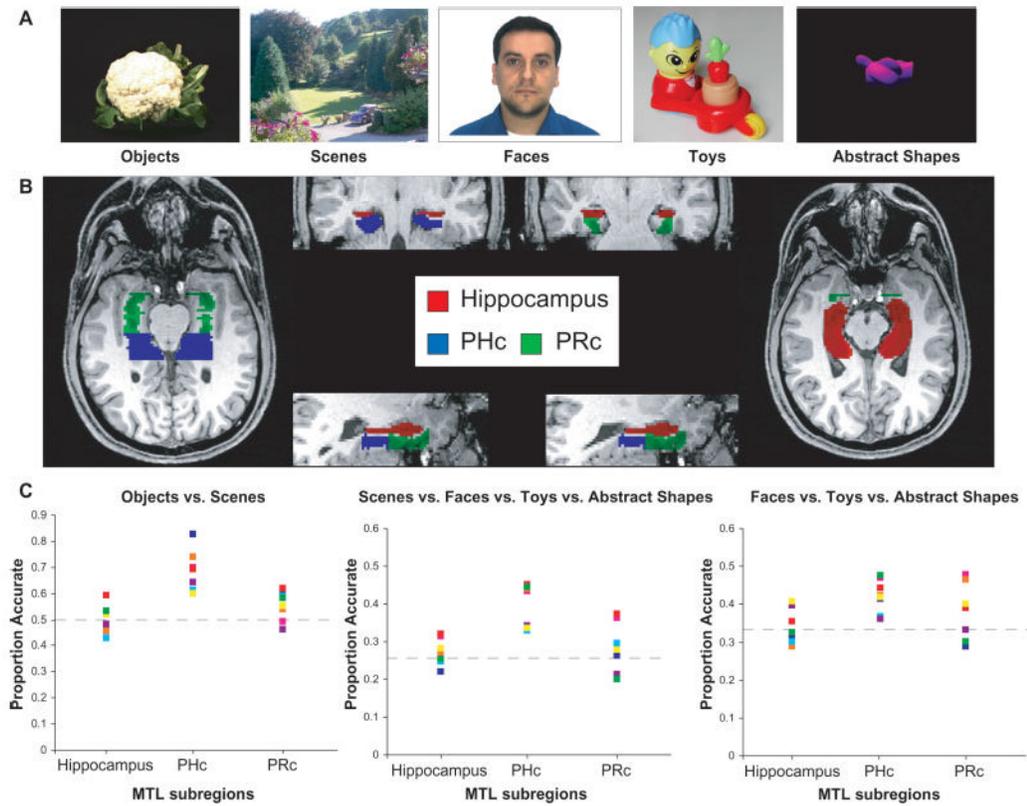
FIGURE 1.
(A) Sample stimuli. (B) Masked regions for a single subject. (C) Individual subject classification results. Chance performance level is indicated by the gray dashed lines (50%, 25%, and 33%, respectively).

TABLE 1

Difference in Activation Between the Correct Output Unit and the Average Activation of the Incorrect Output Units, Providing a Measure of Classification Confidence, for Objects vs. Scenes (A), Scenes vs. Faces vs. Toys vs. Abstract Shapes (B), and Faces vs. Toys vs. Abstract Shapes (C)

	Output unit (A)	
	Object	**Scene**
PHc	0.29	0.42
PPA	0.41	0.45
PRc	0.12	0.09
Hippocampus	−0.01	0.01

	Output unit (B)			
	Scene	Face	Abstract shape	Toy
PHc	0.39	0.15	0.15	0.12
PPA	0.40	0.13	0.12	0.11
PRc	0.03	0.07	0.05	−0.01
Hippocampus	0.02	−0.04	0.04	0.04

	Output unit (C)		
	Face	Abstract shape	Toy
PHc	0.13	0.16	0.10
PPA	0.09	0.11	0.11
PRc	0.07	0.12	0.03
Hippocampus	−0.03	0.07	0.04

TABLE 2

	Mean # voxels in mask	Mean # voxels used in classification	Signal-to-Noise Ratio
Hippocampus	808	267	24.17
PHc	1108	391	25.02
PRc	725	292	23.03
PPA	74	39	21.14

Table shows the mean number of voxels per ROI, the mean number of voxels that were used in the classification analyses (i.e., after using an initial ANOVA to remove uninformative voxels), and the mean signal-to-noise ratio per mask (see Detailed Methods).