

Neural responses to visual scenes reveals inconsistencies between fMRI adaptation and multivoxel pattern analysis

Russell A. Epstein*, Lindsay K. Morgan

Department of Psychology, University of Pennsylvania, 3720 Walnut St., Philadelphia, PA 19104, USA

ARTICLE INFO

Article history:

Received 5 April 2011

Received in revised form

25 September 2011

Accepted 27 September 2011

Available online 5 October 2011

Keywords:

Visual scene recognition

fMRI adaptation

Multivoxel pattern analysis

Parahippocampal cortex

Retrosplenial cortex

Spatial navigation

ABSTRACT

Human observers can recognize real-world visual scenes with great efficiency. Cortical regions such as the parahippocampal place area (PPA) and retrosplenial complex (RSC) have been implicated in scene recognition, but the specific representations supported by these regions are largely unknown. We used functional magnetic resonance imaging adaptation (fMRIa) and multi-voxel pattern analysis (MVPA) to explore this issue, focusing on whether the PPA and RSC represent scenes in terms of general categories, or as specific scenic exemplars. Subjects were scanned while viewing images drawn from 10 outdoor scene categories in two scan runs and images of 10 familiar landmarks from their home college campus in two scan runs. Analyses of multi-voxel patterns revealed that the PPA and RSC encoded both category and landmark information, with a slight advantage for landmark coding in RSC. fMRIa, on the other hand, revealed a very different picture: both PPA and RSC adapted when landmark information was repeated, but category adaptation was only observed in a small subregion of the left PPA. These inconsistencies between the MVPA and fMRIa data suggests that these two techniques interrogate different aspects of the neuronal code. We propose three hypotheses about the mechanisms that might underlie adaptation and multi-voxel signals.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

A central concern of cognitive neuroscience is understanding the information processing functions of different brain regions. A standard approach is to identify the representational distinctions supported by a brain region; that is, which items does a region treat as identical and which does it treat as distinct (and to what extent)? At the neuronal level, such questions are often answered by measuring the tuning curves of single units, or, in more recent treatments, by identifying the distinctions that can be made within multi-unit response spaces (Hung, Kreiman, Poggio, & DiCarlo, 2005). In functional magnetic resonance imaging (fMRI) studies, on the other hand, such questions have been addressed by two techniques: multivoxel pattern analysis (MVPA) and fMRI adaptation (fMRIa). The first approach (MVPA) examines the voxelwise response patterns elicited by different stimuli (or classes of stimuli) to determine which items elicit patterns that are distinguishable (Cox & Savoy, 2003; Haxby et al., 2001; Norman, Polyn, Detre, & Haxby, 2006). The second approach examines the effect of repeating items over time under the hypothesis that repetition of representationally-similar items will elicit a reduced response

(Grill-Spector & Malach, 2001; Grill-Spector, Henson, & Martin, 2006; Kourtzi & Kanwisher, 2001).

Here we use MVPA and fMRIa to understand the neural representations that underlie the recognition of real-world visual scenes. Human observers can analyze the content and significance of scenes quite efficiently (Biederman, 1972; Fei-Fei, Iyer, Koch, & Perona, 2007; Potter, 1975). Brain regions have been identified that respond more strongly to images of real-world scenes (landscapes, cityscapes, rooms) than to images of single objects (vehicles, appliances, animals), bodies or faces (Epstein & Kanwisher, 1998). These include the Parahippocampal Place Area (PPA) and the Retrosplenial Complex (RSC). Although these earlier results, along with concomitant neuropsychological data (Epstein, DeYoe, Press, Rosen, & Kanwisher, 2001; Habib & Sirigu, 1987; Mendez & Chierri, 2003; Takahashi, Kawamura, Shiota, Kasahata, & Hirayama, 1997) suggest that the PPA and RSC play an important role in scene processing, the specific functions that these regions play in scene recognition remain undetermined. In particular, it is unclear whether these regions primarily support identification in terms of general categories (e.g. beach, desert, kitchen, bedroom) or as specific exemplars (e.g. the kitchen on the fifth floor of the Penn Center for Cognitive Neuroscience) (Epstein & Higgins, 2007). Whereas categorical information is important for making predictions about what kind of actions or events are likely to be found in a scene (Bar, 2004), exemplar information is important for spatial

* Corresponding author. Tel.: +1 215 573 3532.

E-mail address: epstein@psych.upenn.edu (R.A. Epstein).

navigation when different places need to be identified and distinguished (Epstein, Parker, & Feiler, 2007).

Recent MVPA studies have made progress on these issues. Walther, Caddigan, Fei-Fei, and Beck (2009) demonstrated that multi-voxel patterns (MVPs) in the PPA and RSC discriminate between six scene categories. Interestingly, above-chance levels of classification performance were observed in the object-selective lateral occipital complex (LOC) and early visual cortex (EVC), regions not generally associated with scene processing (although see MacEvoy & Epstein, 2011; Park, Brady, Greene, & Oliva, 2011). However, multi-voxel patterns in the PPA (and, to a lesser extent, RSC) appeared to have a tighter relationship with recognition performance than MVPs in other brain areas: when MVPA classification errors were compared to errors made by human subjects, both the PPA and human observers tended to get confused about the same category pairs. This finding parallels similar results on object recognition, where object identity can be decoded from MVPs in both LOC and early visual cortex, but only LOC activity patterns predict behavioral performance (Williams, Dang, & Kanwisher, 2007). Walther et al.'s results implicate the PPA in scene categorization, but do not exclude the possibility that it might also be involved in the identification of specific scenes. Indeed, a recent report from our laboratory found that MVPs in the PPA and RSC reliably distinguished between individual landmarks on a familiar college campus (Morgan, Macevoy, Aguirre, & Epstein, 2011). Thus, the PPA and RSC might be involved in both kinds of scene recognition.

These MVPA findings complement earlier studies that investigated PPA and RSC scene representations using fMRIa. These studies found reduced response in the PPA and RSC when individual scenes were repeated, suggesting that these regions encode individual scene exemplars. An important concern of these earlier fMRIa studies was determining the viewpoint-specificity of the repetition effect. An early study using a short-interval repetition paradigm found a purely viewpoint-specific effect: when the first item followed the second item after an interval of only a few hundred msec, adaptation (i.e. reduced response) was observed when the items were identical images, but not when they were images of the same scene taken from different vantage points (Epstein, Graham, & Downing, 2003). Later studies, on the other hand, found some degree of viewpoint tolerance when the first and second item were presented at a much longer repetition interval of several minutes (Epstein, Higgins, & Thompson-Schill, 2005; Epstein, Higgins, Jablonski, & Feiler, 2007). However, even in this case, there was some additional adaptation observed when scenes were repeated from the same view, indicating some degree of viewpoint-specificity even in the face of considerable viewpoint-tolerance. Importantly, both methods revealed adaptation effects that were elicited by *specific scenes*: a place or landmark elicited a reduced response if it had been seen before in the experiment, but not if it was presented for the first time. To our knowledge, adaptation for scene category repetitions has not been previously examined.

As the above discussion indicates, the fMRIa findings on scene processing are not entirely congruent with the MVPA findings. On the one hand, both sets of findings implicate the PPA and RSC in scene recognition – the MVPA results because of the strong relationship between multi-voxel patterns and behavioral distinctions, the fMRIa results because adaptation effects were generally restricted to the PPA, RSC, or a third scene-responsive region in the transverse occipital sulcus. On the other hand, the two sets of findings seem to disagree about the level at which scenes are represented in the PPA and RSC: MVPA results argue for more categorical representations, while the fMRIa results argue for more specific representations that distinguish between individual scenes or even individual views. These incongruencies do not, however, necessarily indicate a fundamental inconsistency. Although both MVPA and fMRIa provide information about representational distinctions, it is unclear how

these distinctions are instantiated at the neuronal level. Thus it is by no means certain that representational distinctions obtained by one technique should correspond to representational distinctions obtained by the other. In fact, incongruencies between MVPA and fMRIa results have been observed previously in the literature (Drucker & Aguirre, 2009) and exploration of these differences can potentially provide insight into the mechanisms that underlie each signal – a theme that we will explore in this paper.

The current study attempted to clarify some of these outstanding issues regarding the neural representations that underlie scene processing in the PPA and RSC. We were especially interested in two questions. First, to what extent do these regions support recognition of scenes at either the categorical or the individual exemplar level? Second, to what extent do MVPA and fMRIa give consistent results? To address these questions, we scanned subjects with fMRI while they viewed images drawn from 10 outdoor categories and 10 familiar landmarks from the Penn campus. Stimuli were presented in a continuous carryover design, which counterbalances main effects and carry-over effects, thus allowing MVPs and fMRIa to be analyzed in the same data set (Aguirre, 2007). We have previously presented some of the data from the Penn landmarks (Morgan et al., 2011), but the Outdoor Category data, along with most of the analyses, are new.

To anticipate, our results suggest that the PPA might support recognition of scenes at both the categorical and individual exemplar level while RSC might be more involved in recognition of specific familiar places. Furthermore, our data indicate some striking dissociations between the representational distinctions revealed by MVPA and the representational distinctions revealed by fMRIa, which suggests that these techniques index fundamentally different aspects of the neural code.

2. Materials and methods

2.1. Subjects

Fifteen healthy, right-handed volunteers (10 female; mean age, 22.6 years) with normal or corrected-to-normal vision were recruited from the University of Pennsylvania community. All subjects gave written informed consent according to procedures approved by the University of Pennsylvania institutional review board.

2.2. Stimuli and procedure

Stimuli were digitized color photographs of 10 outdoor scene categories (e.g., beach, playground) and 10 prominent landmarks (i.e., buildings and statues) from the University of Pennsylvania campus (Fig. 1). The Penn landmarks were familiar to all subjects; the outdoor category images depicted unfamiliar locations. Outdoor categories were chosen to be roughly equivalent to “basic level” scene categories identified by Tversky and Hemenway (1983) as being the preferred level of description for scenes; these categories tend to have characteristic objects and perceptual features (e.g., sand, water, and palm trees for beach) and are associated with certain activities that are appropriate to that setting (e.g. swimming, sunbathing). Penn landmarks were chosen to be prominent fixed environmental items whose identity and location were familiar to most Penn students. We obtained 22 distinct exemplar photographs (e.g. 22 different beaches) for each category and 22 distinct views of each landmark for a total of 440 images in all (for examples, see Supplementary Figure). Images were presented at 1024 × 768 pixel resolution and subtended a visual angle of 22.9° × 17.4°.

All 440 images were presented without repetition over the course of 4 fMRI scans that lasted 6 min 51 s each. In counterbalanced order, subjects viewed 2 runs of outdoor scene categories and 2 runs of campus landmarks (i.e., scene categories and campus landmarks never appeared within the same run). Images were presented every 3 s in a continuous-carryover sequence that included 6 s null trials interspersed with the stimulus trials (Aguirre, 2007). This stimulus sequence counterbalances main effects and first-order carry-over effects by ensuring that each category (or landmark) is preceded by every other category (or landmark) equally often. This counterbalancing ensures independence between the main effects (used for MVPA) and the first-order carryover effects (used to assess adaptation), thus allowing one to use the same fMRI dataset for both analyses. Two unique continuous-carryover sequences were defined for each subject (one for the category runs; the other for the landmark runs). On each stimulus trial, an image of a scene category or landmark was presented for 1 s followed by 2 s of a grey screen with a black fixation cross. Subjects were asked to covertly identify the scene category or campus

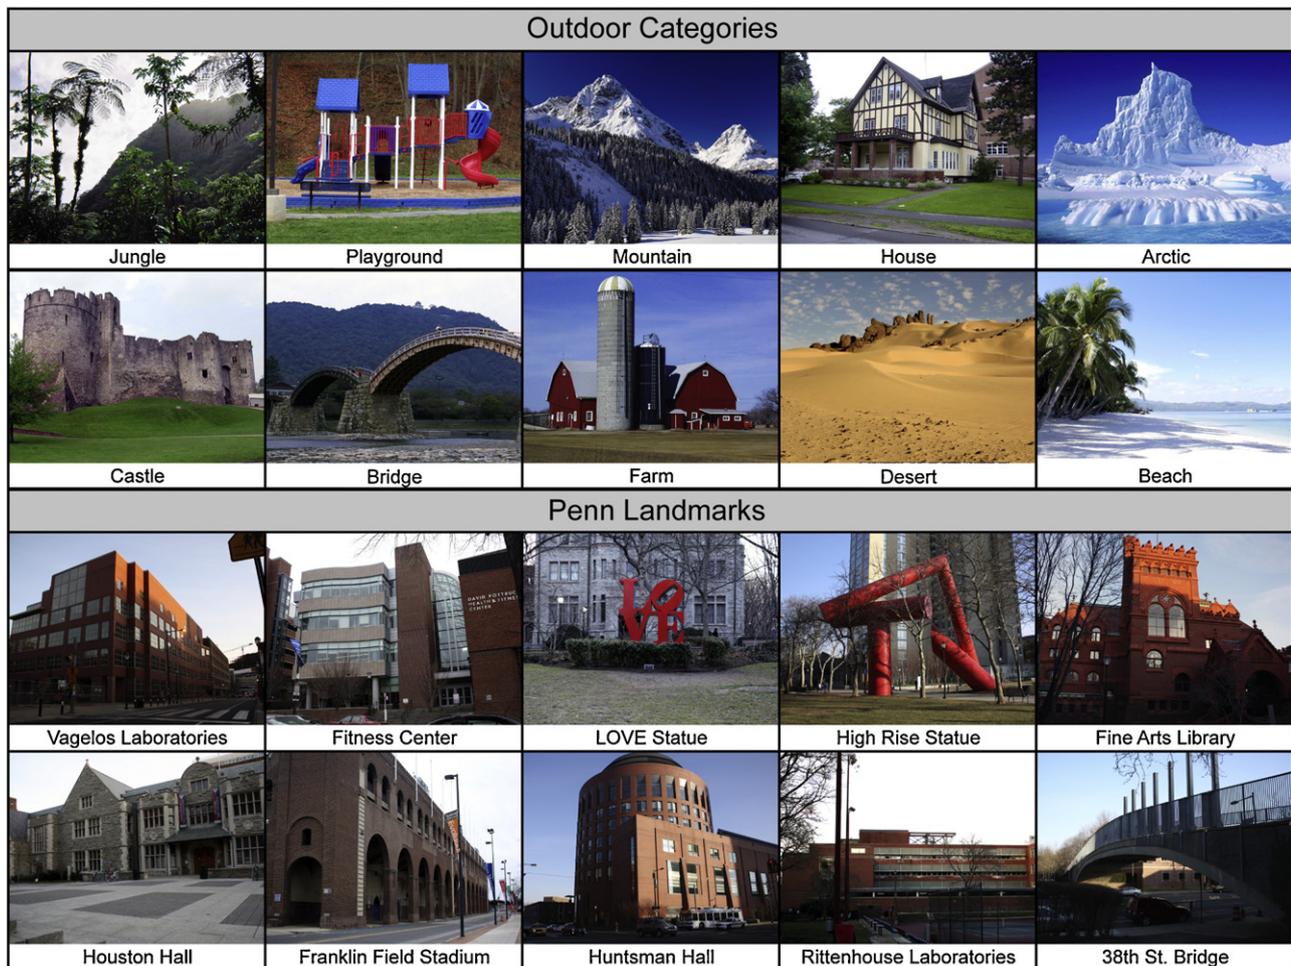


Fig. 1. Examples of the 10 outdoor categories and 10 Penn landmarks displayed during the experiment. 22 different images were shown for each category and landmark (for examples see [Supplementary Figure](#)).

landmark and make a button press once they had done so. During null trials, a grey screen with black fixation cross was presented for 6 s during which subjects made no response.

After the main experiment, 2 functional localizer scans were administered. Subjects performed a one-back repetition task while they viewed 18-s blocks of images of places (e.g., cityscapes, landscapes), single objects without backgrounds, grid-scrambled objects, and other stimuli, presented for 490 ms with a 490 ms inter-stimulus interval. Each scan lasted 7 min 48 s.

2.3. fMRI acquisition

Scans were performed at the Hospital of the University of Pennsylvania on a 3T Siemens Trio scanner equipped with a Siemens body coil and an eight-channel head coil. High resolution T1-weighted anatomical images were acquired using a 3D MPRAGE pulse sequence (TR = 1620 ms, TE = 3 ms, TI = 950 ms, voxel size = $0.9766 \times 0.9766 \times 1$ mm, matrix size = $192 \times 256 \times 160$). T2*-weighted images sensitive to blood oxygenation level-dependent (BOLD) contrasts were acquired using a gradient-echo echo-planar pulse sequence (TR = 3000 ms, TE = 30 ms, flip angle = 90° , voxel size = $3 \times 3 \times 3$ mm, matrix size = 64×64 , 46 axial slices).

2.4. fMRI data analyses

2.4.1. Preprocessing

Prior to analysis, functional images were corrected for differences in slice timing by resampling slices in time to match the first slice of each volume, realigned to the first image of the scan, and spatially normalized to the Montreal Neurological Institute template using a linear 12-parameter affine transformation as implemented in SPM2. MR values for each scan run were mean scaled to 1 prior to analysis to ensure that beta weights extracted using the general linear model corresponded to percent signal change. Data used for the region of interest definition and fMRI adaptation analyses were spatially smoothed with a 6-mm FWHM Gaussian filter; data for all other analyses were left unsmoothed. Analyses of fMRI timecourses were performed using the general linear model as implemented in VoxBo ([www.voxbo.org](#)),

including an empirically-derived $1/f$ noise model, filters that removed high and low temporal frequencies, regressors to account for global signal variations, and regressors to account for differences in the mean level of activation between scan runs.

2.4.2. Regions of interest

Data from the functional localizer scans were used to define several regions of interest (ROIs) in each subject based on preferential response to scenes, objects, or low-level visual features (Fig. 8c). The PPA and RSC were defined as the set of voxels in the collateral sulcus/posterior parahippocampal region (PPA) or retrosplenial/medial parietal region (RSC) that responded more strongly to scenes than to objects. We also identified a third scene-responsive region in the transverse occipital sulcus (TOS) using the same contrast. The lateral occipital complex (LOC) was defined as the region of lateral/ventral occipitotemporal cortex that responded more strongly to objects than to scrambled objects. Early visual cortex (EVC) was defined as the region extending from the occipital pole that responded more strongly to grid-scrambled objects than to intact objects. Thresholds were determined on a subject-by-subject basis to be consistent with those identified in previous studies and ranged from $T > 2.0$ to $T > 3.5$ (mean $T = 2.7$). Bilateral PPA and LOC were located in all 15 subjects. Right RSC was identified in all subjects, left RSC in 13/15 subjects, EVC (not differentiated into hemisphere) in 14/15 subjects, and both left and right TOS in 13/15 subjects.

We further divided each subject's PPA along the collateral sulcus in each hemisphere to create 4 subregions (left lateral, left medial, right lateral, right medial). This was done using ITK-SNAP ([www.itksnap.org](#)) in the following manner. First, the collateral sulcus was identified in the coronal plane on the most posterior slice of the PPA. Next, the sulcus was traced from the fundus to the cortical surface and the visibility of the PPA was toggled on. The plane of the sulcus was elongated if necessary to capture the entire extent of the PPA. Finally, the PPA visibility was toggled off and parcellation proceeded to the next anterior slice. If multiple branches of the collateral sulcus were present on any slice, the main branch was identified in the sagittal view.

2.4.3. Classification from multivoxel patterns

To determine whether multi-voxel patterns within each ROI encoded information about the scene category or landmark being viewed, we implemented a standard classification technique in which multi-voxel patterns were compared across scan runs (Haxby et al., 2001). Outdoor category runs were analyzed separately from landmark runs. In both cases, we used a general linear model to estimate the magnitude of the response at each voxel for the 10 categories (or landmarks) in each scan run. Specifically, each GLM consisted of 20 regressors (10 conditions \times 2 scan runs) in which each stimulus presentation event was modeled as a unit impulse function convolved with a canonical hemodynamic response function. Beta values (corresponding to percent signal change) for each of the 20 regressors in the model were then extracted at each voxel. Classification was performed on these beta values using the method of pairwise comparison described by Haxby et al. (2001). A cocktail mean pattern consisting of the average response across all scene categories (or landmarks) was calculated for each scan run and subtracted from the individual patterns; the patterns for all 10 categories (or landmarks) were then compared across scan runs using Euclidean distance as a measure of similarity between conditions. Patterns were considered correctly classified if within-condition distances (e.g., Beach-Beach) were smaller than between-condition distances (e.g., Beach-Playground). Classification accuracy was averaged across all possible pairwise comparisons for a given ROI and tested against random chance (i.e., 0.5) using a one-tailed *t*-test. Classification performance was substantially unchanged when correlation rather than Euclidean distance was used to evaluate similarities between activation patterns.

2.4.4. Gamut analysis

To test for a difference in the gamut of outdoor category representations and Penn landmark representations, we computed Euclidean distances between multivoxel patterns for each category–category and landmark–landmark pairing for both PPA and RSC. These Euclidean distances were the same values that were used for the MVPA classification analysis. However, in this analysis we did not perform the additional step of comparing Euclidean distances between pairings, as this step eliminates information about the absolute distances between response vectors. Rather, we simply averaged Euclidean distances across all within-category, between-category, within-landmark, and between-landmark pairings. We then compared these values across stimulus classes (i.e. categories or landmarks) to determine if the response vectors for either of the two stimulus classes covered a larger portion of the response space.

2.4.5. Comparison of neural and visual dissimilarity

To test the hypothesis that multi-voxel patterns might reflect coding of visual properties, we computed the visual dissimilarity between the scene categories and between the Penn landmarks using a texture model that has previously been shown to perform similarly to human subjects performing scene identification on brief image presentations (<70 ms) (Renninger & Malik, 2004). Images were converted to grayscale and passed with V1-like filters to generate a list of the 100 most prototypical texture features found across the images (MATLAB code available at renningerlab.org). A histogram of texture frequency was then generated for each image. The visual dissimilarity between a pair of images was calculated by comparing the distribution of the two histograms using a χ^2 measure (smaller χ^2 values correspond to more similar images). We computed the dissimilarity for every pair of images both within a category/landmark (e.g., Beach1 vs. Beach2) as well as across categories/landmarks (e.g., Beach1 vs. Playground1). We then generated visual confusion matrices that represented the average visual dissimilarities between categories or landmarks by averaging over all the relevant pairwise dissimilarities.

Complementary neural confusion matrices were created by calculating Euclidean distances between multivoxel patterns across runs. These distances were calculated for each landmark–landmark and each category–category pairing, and were the same values used for the classification and gamut analyses. Neural confusion matrices were generated for each ROI in each subject for both scene categories and landmarks. We next created average neural confusion matrices for each ROI by averaging across subjects at each cell of the confusion matrix. Finally, to test whether neural similarity in the ROIs was related to low level visual similarity, we computed the correlation between the off-diagonal elements of the average neural confusion matrix and the visual confusion matrix using Pearson's *R*.

2.4.6. fMRI adaptation (fMRIa)

Because we presented stimuli in a continuous carry-over sequence that fully counterbalances main effects and first order carry-over effects, we were able to interrogate both multivoxel patterns and fMRIa using the same dataset. In particular, we examined the effect of repeating images that belong to the same scene category (e.g., Beach Image 1 \rightarrow Beach Image 2) and the effect of repeating images of the same campus landmark (e.g., Houston Hall Image 1 \rightarrow Houston Hall Image 2). To this end, we created a general linear model that contained regressors for (1) repetition of a category/landmark, (2) response to any visual stimulus versus baseline and (3) trials that appear immediately after null trials. Beta values for the repetition regressor were extracted at each ROI for every subject, converted to percent signal change, averaged across subjects, and compared to zero using a one-tailed *t*-test.

This analysis was performed separately for the outdoor categories and for the Penn landmarks.

2.4.7. Voxelwise informativeness

To examine the consistency of a region's pattern of responses to categories and landmarks, we calculated voxelwise "informativeness" measures as follows. For each voxel in the ROI, we extracted an activity vector corresponding to the response to all 10 categories in the first category scan run, a separate activity vector corresponding to the response to all 10 categories in the second category scan run, and then calculated the correlation between these vectors. We then repeated the procedure for the landmark runs. This measure expresses the potential informativeness of the voxelwise response for activity-based analyses, because voxels that display reliable response patterns across scan runs will have high between-run correlation values (Kravitz, Peng, & Baker, 2010; Mitchell et al., 2008). We use these informativeness values in two ways. First, we averaged the values for all the voxels contained within an ROI, in order to determine the average informativeness for each region. Second, we examined the relationship between informativeness and fMRIa by correlating informativeness values with the beta value of the repetition effect across all voxels in an ROI. For this analysis, repetition beta values were calculated from unsmoothed fMRI data. We then transformed each subject's Pearson's *R* value to a Fisher *z* value, averaged across subjects, and tested whether the average correlation was different from zero using a two-tailed *t*-test.

2.4.8. Whole-brain analyses

Although our focus was on response within the predefined ROIs, we also performed exploratory whole-brain analyses to determine whether any areas outside our ROIs exhibited above-chance MVPA classification performance or significant fMRI adaptation. For the whole-brain version of MVPA, we implemented the searchlight method of Kriegeskorte, Goebel, and Bandettini (2006), which obtains a measure of classification performance in the neighborhood surrounding each voxel of each subject's brain. Specifically, we defined a small spherical ROI (the "searchlight"; radius = 5 mm) centered on each voxel and used the pairwise comparison method described earlier to calculate classification performance for outdoor categories and, separately, for Penn landmarks within the sphere; these classification values were then assigned to the central voxel. For the whole-brain analysis of fMRIa, we used the general linear model described previously to calculate beta values for category repetition and landmark repetition at each brain voxel. Subject-specific searchlight and adaptation maps were entered into a second-level random effects analysis to determine whether values were significantly greater than chance (for searchlight analyses) or significantly different than zero (for adaptation analyses). Both searchlight and adaptation maps were smoothed to 9-mm FWHM before entry into the second-level analysis. Monte-Carlo simulations involving sign permutations of the whole-brain data from individual subjects (1000 relabelings, 12 mm FWHM pseudo-*t* smoothing) were performed to find the true Type I error rate (Nichols & Holmes, 2002), thus correcting for the fact that statistical comparisons were made simultaneously for all voxels in the brain.

3. Results

3.1. Decoding landmarks and outdoor categories with MVPA

A primary goal of this study was to understand how PPA and RSC might encode scene categories and individual scene exemplars. As a first step, we used standard MVPA techniques to verify that these regions distinguish between scenes at both of these representational levels. Classification performance (Fig. 2) was well above

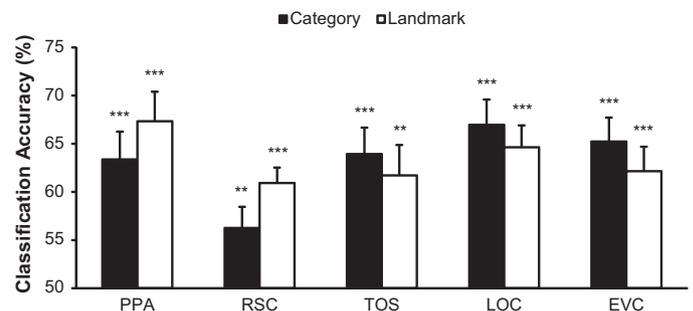


Fig. 2. MVPA classification accuracy (mean \pm SEM) in functional ROIs. Both outdoor categories and Penn landmarks could be decoded at above-chance (>50%) levels in all regions. PPA, parahippocampal place area; RSC, retrosplenial complex; TOS, transverse occipital sulcus; LOC, lateral occipital complex; EVC, early visual cortex; ***p* < 0.01; ****p* < 0.001.

chance for outdoor scene categories [PPA $t(14)=4.6$, $p=0.0002$; RSC $t(14)=2.8$, $p=0.007$] and also for individual Penn landmarks [PPA $t(14)=5.6$, $p=0.00003$; RSC $t(14)=6.8$, $p=0.000004$] consistent with previous results (Morgan et al., 2011; Walther et al., 2009). When classification performance for outdoor categories was directly compared to classification performance for individual landmarks, there was no difference in the PPA [$t < 1$, n.s.], but classification performance was higher for landmarks than for categories in RSC [$t(14)=2.2$, $p=0.04$, two-tailed]. Thus, multi-voxel patterns in the PPA and RSC convey information about both the category and specific identity of a scene, at about the same level of accuracy in the PPA, but with greater accuracy for identity than for category in RSC.

These results were not restricted to the PPA and RSC. We could decode scene category with a high degree of accuracy in the lateral occipital complex [LOC, $t(14)=6.5$, $p=0.000007$], early visual cortex [EVC, $t(13)=6.2$, $p=0.00002$], and transverse occipital sulcus [TOS, $t(12)=5.1$, $p=0.0001$]. Similarly high levels of accuracy were obtained for decoding of Penn landmarks in all three regions [LOC, $t(14)=6.4$, $p=0.000008$; EVC $t(13)=4.8$, $p=0.0002$; TOS $t(12)=3.7$, $p=0.0015$]. These results are consistent with earlier studies demonstrating decoding of high-level scene categories in these regions, findings that are likely reflective of reliable differences in diagnostic objects and shapes (for LOC), low-level visual properties (for EVC), and low-level scene properties (for TOS). Although classification performance was numerically higher for outdoor categories than for Penn landmarks in all three regions, these differences in performance were not significant (all $ps > 0.4$). Here we focus our attention primarily on the PPA and RSC, as previous work has suggested that multi-voxel codes in these regions are most closely tied to scene recognition performance (Walther et al., 2009).

To test whether category and landmark information might be restricted to certain subregions of the PPA and RSC, we examined response for each hemisphere separately. We also further subdivided the PPA into territory lateral and medial to the collateral sulcus (Fig. 3), as previous work suggests that PPA might consist of two subregions (Arcaro, McMains, Singer, & Kastner, 2009) for which the collateral sulcus is a plausible anatomical boundary (Sewards, 2010). In the PPA, classification of outdoor scene categories was significantly above chance in three of the four subregions [left lateral $t(13)=2.8$, $p=0.007$; left medial $t(14)=4.4$, $p=0.0003$; right lateral $t(14)=4.9$, $p=0.0001$] with the only exception being the right medial PPA [$t < 1$, n.s.]. Classification of Penn landmarks was above chance in all four subregions [left lateral $t(13)=4.8$, $p=0.0002$; left medial $t(14)=3.0$, $p=0.005$; right lateral $t(14)=5.3$, $p=0.00006$; right medial $t(13)=3.9$, $p=0.001$]. In RSC, classification of outdoor categories was above chance in both hemispheres [left $t(12)=3.0$, $p=0.005$; right $t(14)=1.9$, $p=0.04$] as was

classification of Penn landmarks [left $t(12)=6.2$, $p=0.00002$; right $t(14)=6.2$, $p=0.00001$].

3.2. Gamut analysis

In the MVPA analyses above, classification was based on comparison of within-category/landmark neural dissimilarities to between-category/landmark neural dissimilarities, where neural dissimilarity was defined by Euclidean distances between multi-voxel patterns. Such an approach is standard in MVPA. We hypothesized that this approach could potentially obscure differences between the neural codes supporting the coding of the two stimulus classes. In particular, because the MVPA classification scheme involves comparing Euclidean distances *within* each stimulus class, rather than *across* stimulus classes, it might obscure between-class differences in the underlying representational spaces.

We were especially concerned with this issue for the following reason. Even a cursory examination of the stimulus set makes it evident that the outdoor category images are more visually disparate than the Penn landmark images (for examples, see Fig. 1 and Supplementary Figure). Furthermore, the outdoor categories might be considered to be more semantically disparate, given that the ten Penn landmarks can be grouped into fewer than ten categorical descriptors (Building, Statue, Stadium, Bridge). Given these differences, it is somewhat surprising that classification performance is equivalent for both outdoor categories and Penn landmarks in the PPA (and, indeed, better for the Penn landmarks in RSC).

One possibility is that patterns corresponding to the ten Penn landmarks might be more similar to each other but also more reliable across scan runs than the patterns corresponding to the ten outdoor categories. For example, beaches and jungles might elicit neural patterns that are rather dissimilar while Huntsman Hall and Houston Hall might elicit neural patterns that are rather similar; but at the same time, beach and jungle patterns might vary considerably across runs while Huntsman Hall and Houston Hall patterns might be more consistent. One might, therefore, get equivalent classification performance for Penn landmarks and outdoor categories despite widely different gamuts for these two disparate stimulus classes.

To test this idea, we simply plotted the average Euclidean distance for within-category/landmark and between-category/landmark pairs, separately for the outdoor categories and the Penn landmarks (Fig. 4). A 2×2 ANOVA revealed that within-category and within-landmark distances were significantly smaller than between-category and between-landmark distances in both the PPA [$F(1, 14)=56.4$, $p=0.000003$] and RSC [$F(1, 14)=25.4$, $p=0.0002$], as one would expect given the above chance

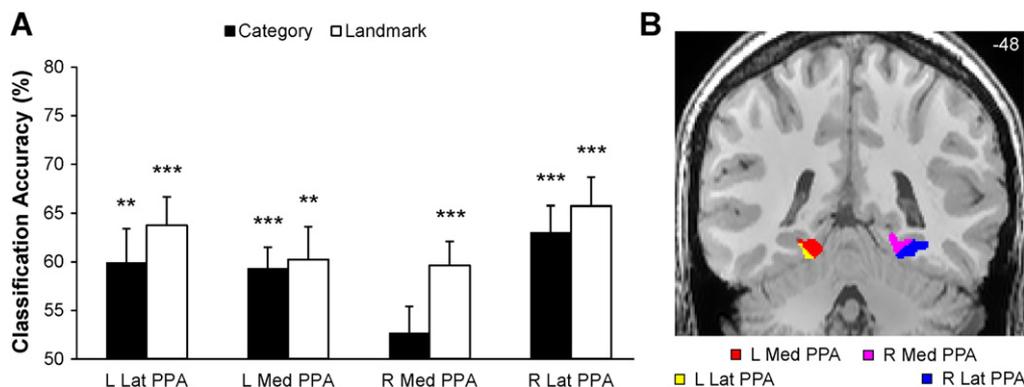


Fig. 3. MVPA classification accuracy within PPA subregions. (A) Classification accuracy was significantly above chance (>50%), or nearly so, for both stimulus classes in most PPA subregions. Numbers are mean \pm SEM. (B) An example of the anatomical locations of the 4 PPA subregions for one coronal slice from 1 subject. The lateral/medial boundary is the collateral sulcus. Lat, lateral; Med, medial; ** $p < 0.01$; *** $p < 0.001$.

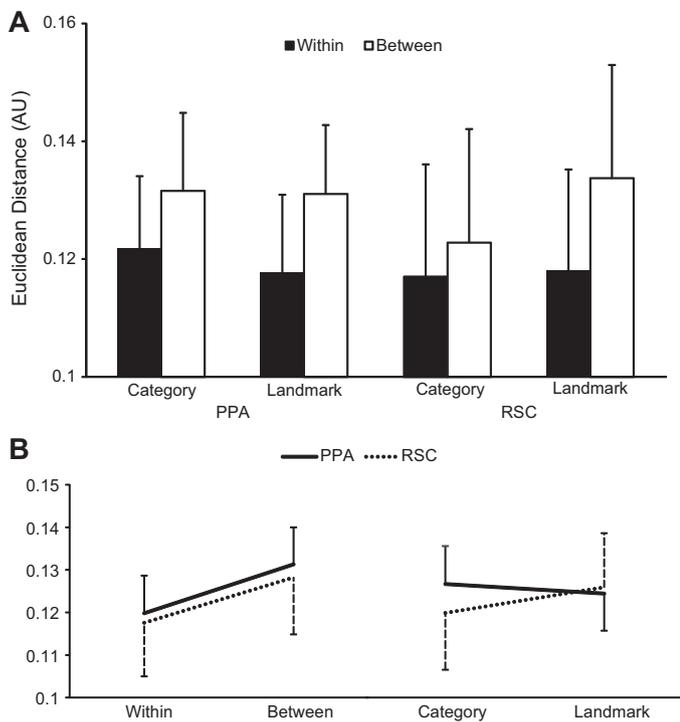


Fig. 4. Gamut analysis in PPA and RSC. (A) Average Euclidean distances (mean \pm SEM) between multivoxel response patterns evoked in different scan runs. These distances were calculated for each category–category and landmark–landmark pairing and then averaged separately across all same-category/landmark pairs (*within* pairings) and across all different-category/landmark pairs (*between* pairings). AU, arbitrary units of Euclidean distance in fMRI response space. (B) Euclidean distances were greater for between-category/landmark pairings than for within-category/landmark pairings in both regions (left panel). Although the main effect of category vs. landmark was not significant (right panel) in either region, there was a significant stimulus class (category vs. landmark) by region interaction, whereby RSC showed relatively larger gamut for Penn landmarks, while PPA showed relatively larger gamut for outdoor categories.

classification performance (Fig. 4b, left panel). Contrary to our hypothesis, however, average Euclidean distances between patterns were equivalent for the Penn landmarks and outdoor categories in the PPA [$F < 1$, n.s.]. In RSC, there was a non-significant trend towards a larger gamut for Penn landmarks than for outdoor categories [$F(1, 14) = 2.7$, $p = 0.12$] along with a significant interaction between stimulus class and type of pairing [$F(1, 14) = 7.62$, $p = 0.015$], reflecting the fact that within vs. between category differences were larger for the Penn landmarks than for the outdoor categories in this region (again, consistent with the previous classification results).

Although these results do not support the hypothesis that the gamuts differ between the Penn landmarks and the outdoor categories in the PPA, they do emphasize some intriguing differences between the PPA and the RSC. Most notably, although the main effect of outdoor category vs. Penn landmark was not significant in either region, the two nonsignificant trends ran in opposite directions (Fig. 4b, right panel). That is, whereas PPA had a very weak tendency to consider the outdoor categories to be more disparate than the Penn landmarks, RSC treated the Penn landmarks as the more representationally disparate stimulus class. Indeed, when the data from the two ROIs were combined into a single ANOVA, there was a significant interaction of ROI with stimulus class [$F(1, 14) = 6.7$, $p = 0.02$]. Furthermore, between-landmark distances were larger than between-category distances in RSC [$t(14) = 3.7$, $p = 0.003$] but were equivalent in the PPA [$t < 1$, n.s.]. These data suggest that RSC neural codes might be more useful for

distinguishing between different familiar landmarks than for distinguishing between different scene categories (an effect that was also indicated by superior landmark classification in Section 3.1). PPA neural codes, on the other hand, might be equally useful for both scene recognition tasks.

3.3. Relating neural dissimilarities to visual dissimilarities

The previous results would seem to argue against the idea that scenes are coded in the PPA in terms of visual properties, because they failed to find a difference between neural coding of Penn landmarks (which are more visually similar to each other) and outdoor categories (which are more visually dissimilar). Here we perform a more direct test of this idea by examining the relationship between multi-voxel patterns and visual dissimilarity.

To determine visual dissimilarity, we analyzed our stimuli using a texture model that has previously been shown to perform similarly to human subjects tested on scene identification at very brief image presentations (<70 ms) (Renninger & Malik, 2004). Although we make no claims that this algorithm computes quantities that directly correspond to the representations implemented by the PPA or RSC, it provides a reasonable estimate of some of the visual distinctions calculated early in the visual stream. The average visual dissimilarity for between-category image pairs (e.g., farm–beach) was greater than the average visual dissimilarity for within-category image pairs (e.g. farm–farm) [between: 0.51, within: 0.46, $t_{\text{eff}}(24) = 2.3$, $p = 0.032$] and the average visual dissimilarity for between-landmark image pairs was greater than the average visual dissimilarity for within-landmark image pairs [between: 0.33, within: 0.28, $t_{\text{eff}}(13) = 3.7$, $p = 0.003$]. Thus, it is conceivable that classification performance could relate to these visual differences between categories and landmarks.

To test whether this was the case, we constructed the visual confusion matrix corresponding to visual dissimilarities between all outdoor categories and compared it to the equivalent neural confusion matrix; we then did the same analysis for the Penn landmarks. Neural confusion matrices (Fig. 5a) were constructed by calculating the Euclidean distances between multivoxel patterns elicited by outdoor categories or landmarks in different scan runs (i.e. the same quantities that were used for the classification and gamut analyses in the preceding sections). We then plotted neural dissimilarity against visual dissimilarity in each ROI, for both outdoor categories and Penn landmarks (Fig. 5b). Because we expect within-category/landmark dissimilarities to be smaller than between-category/landmark dissimilarities in both the visual and neural domains, we focus on between-category/landmark dissimilarities (i.e. the off-diagonal elements of the confusion matrices).

We did not find any significant relationship between visual and neural dissimilarities for outdoor categories in the scene-responsive ROIs [PPA $r(43) = 0.15$, $p = 0.34$; RSC $r(43) = 0.15$, $p = 0.31$; TOS $r(43) = 0.07$, $p = 0.66$]. In contrast, visual dissimilarity was highly predictive of neural dissimilarity in LOC [$r(43) = 0.39$, $p = 0.008$] and EVC [$r(43) = 0.56$, $p = 0.00007$]. There was no relationship between visual and neural dissimilarity for the Penn landmarks in any of the ROIs [PPA $r(43) = -0.11$, $p = 0.47$; RSC $r(43) = -0.08$, $p = 0.60$; TOS $r(43) = -0.11$, $p = 0.47$; LOC $r(43) = -0.10$, $p = 0.49$; EVC $r(43) = 0.0005$, $p = 0.99$], possibly because the landmarks did not vary sufficiently in the visual domain for a visual–neural correlation to be significant.

In sum, these results suggest that PPA, RSC and TOS do not represent scenes in terms of low-level visual properties. They leave open the possibility they might encode higher-level visual (Greene & Oliva, 2009) or geometric (Kravitz et al., 2010; Park et al., 2011) features—a hypothesis that might be tested by comparing neural space to more sophisticated stimulus feature spaces. Alternatively,

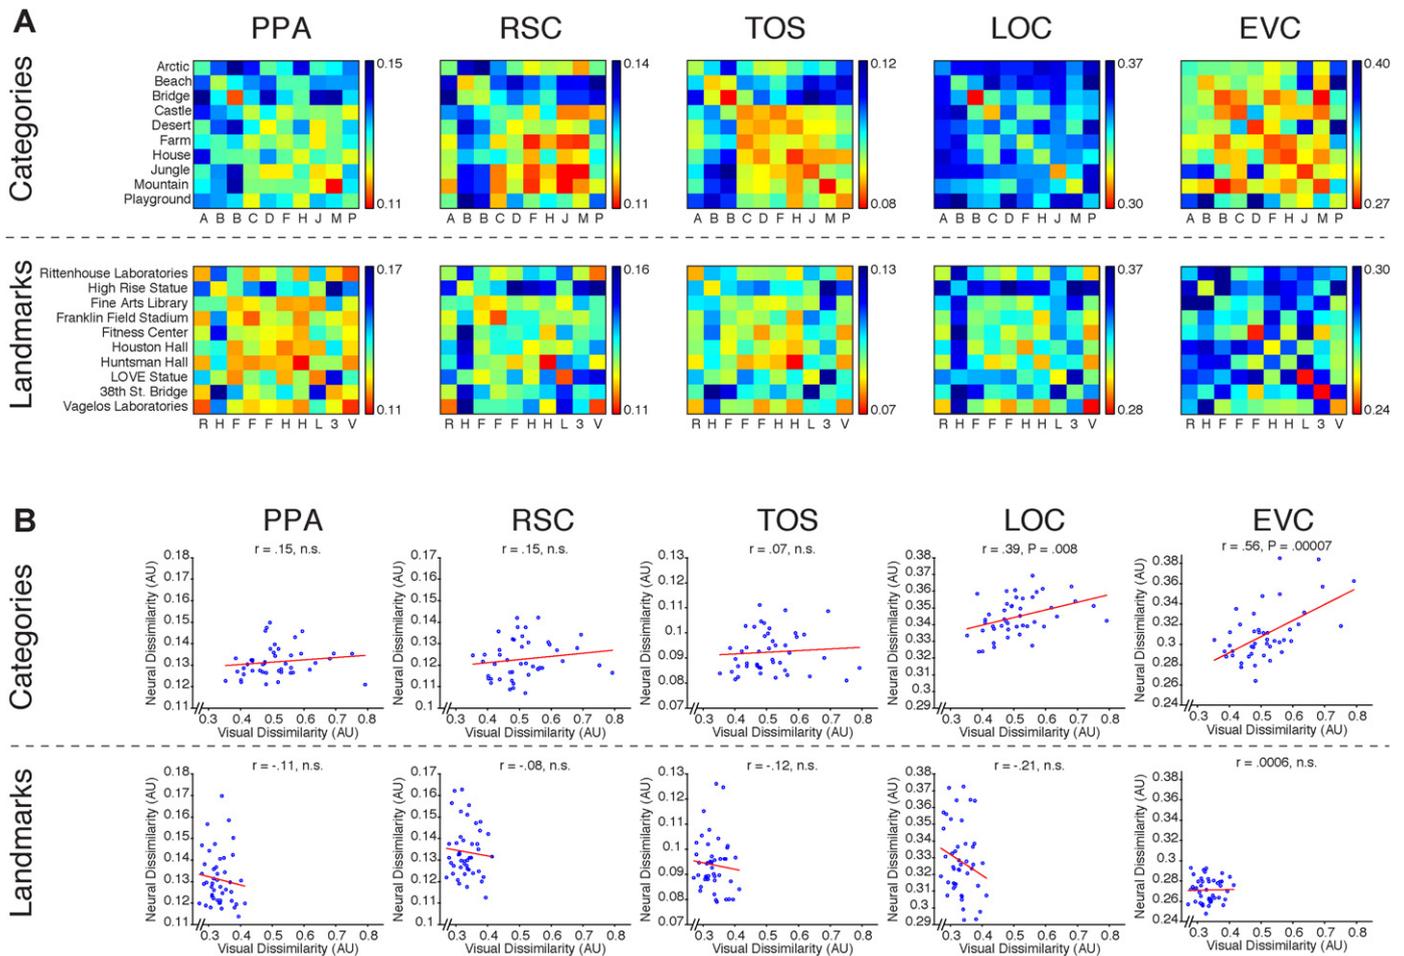


Fig. 5. Comparison of visual vs. neural dissimilarity. (A) Confusion matrices showing neural dissimilarity, defined as Euclidean distance between multivoxel response patterns evoked by the 10 outdoor categories (top row) and the 10 Penn landmarks (bottom row) in different scan runs. Warmer colors indicate more similar patterns (i.e. smaller Euclidean distances) while cooler colors indicate less similar patterns (i.e. larger Euclidean distances). Diagonal elements reflect same-category/landmark pairings; off-diagonal elements reflect different-category/landmark pairings. (B) Neural dissimilarity plotted against visual dissimilarity for each ROI. Each data point represents one category–category or landmark–landmark pairing (for off-diagonal elements of the confusion matrix only). Visual dissimilarity predicts neural dissimilarity for outdoor categories in EVC and LOC, but not PPA, RSC, or TOS. No relationship was observed between visual and neural dissimilarity for Penn landmarks. Note that the range of visual dissimilarities was smaller for the Penn landmarks than for the outdoor categories.

these regions might encode scene categories and landmarks as distinct items independent of their physical features (Walther et al., 2009). In contrast to these null results in scene-responsive regions, both LOC and EVC showed a significant relationship between visual and neural dissimilarity for the outdoor categories, suggesting that these regions might encode low-level visual properties, or object-based features that correlate with low-level visual properties.

3.4. fMRI adaptation effects

In addition to MVPA effects, we also examined fMRI adaptation (fMRIa) effects caused by repetition of category or landmark in successive trials. We were able to look at fMRIa and MVPA simultaneously because we employed a continuous-carryover design that ensured that each outdoor category (or Penn landmark) was preceded equally often by every other outdoor category (or Penn landmark). Thus, for example, beaches were preceded equally often by jungles, farms, castles, deserts, arctic scenes, bridges, and all other outdoor categories including other beaches. This counterbalancing ensured that main effects examined in MVPA and first-order carry-over effects examined in fMRIa were independent of each other. We focus on reductions in fMRI response engendered by repetition of scene category or landmark on successive trials

(beach → beach, Houston Hall → Houston Hall) compared to the “baseline” situation in which category or landmark is not repeated (beach → jungle, Huntsman Hall → Houston Hall).

As a first step, we looked at the effect of repetition on the behavioral response. For each trial, subjects were asked to name the item covertly and press a button once they had done so. We observed behavioral priming effects in both the outdoor category runs (repeat 482 ms, nonrepeat 510 ms, $t(14) = -2.7$, $p = 0.009$) and the Penn landmark runs (repeat 522 ms, nonrepeat 548 ms, $t(14) = -2.0$, $p = 0.03$). That is, responses were speeded when outdoor categories images were preceded by images from the same category, and also when Penn landmark images were preceded by images of the same landmark.

We then looked for an analogous effect on the fMRI response (Fig. 6a). We found a significant reduction of response when Penn landmarks were repeated in PPA [$t(14) = -2.9$, $p = 0.006$], RSC [$t(14) = -3.1$, $p = 0.004$], and TOS [$t(13) = -4.4$, $p = 0.0005$] but only nonsignificant trends in LOC [$t(14) = -1.3$, $p = 0.10$] and EVC [$t(13) = -1.4$, $p = 0.10$]. These findings are generally consistent with previous work indicating that fMRIa effects are found in a more restricted set of regions than MVPA effects; in particular, landmark repetition effects were found in regions that respond preferentially to scenes, but not ROIs that respond preferentially to objects or low-level visual features.

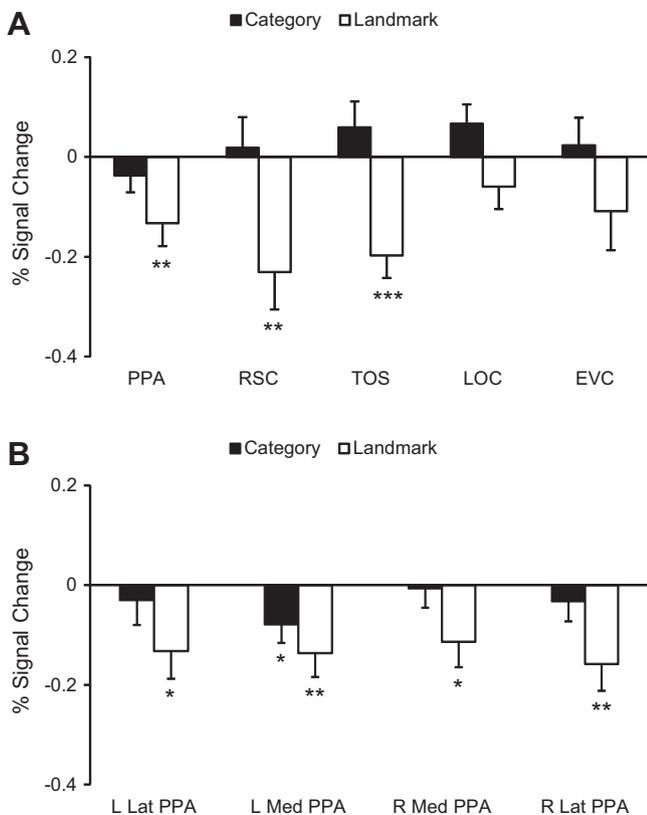


Fig. 6. fMRI adaptation (mean \pm SEM) for category and landmark repetitions. (A) Scene-responsive ROIs (PPA, RSC, TOS) showed adaptation when landmarks were repeated but not when scene categories were repeated. LOC and EVC showed no adaptation for either stimulus class. * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$. (B) Within the PPA, landmark repetition led to adaptation all subregions, whereas category repetition only led to adaptation in the left medial subregion.

In contrast to these robust fMRIa effects for landmarks, we did not observe a reduction of response when outdoor category was repeated in PPA [$t(14) = -1.09, p = 0.15$], RSC [$t < 1, n.s.$], TOS [$t < 1.15, n.s.$] or EVC [$t < 1, n.s.$]. However, a breakdown of the PPA into subregions (Fig. 6b) revealed significant adaptation for category in the left medial portion [$t(14) = -2.1, p = 0.03$; all other subregions $t < 1, n.s.$]. Surprisingly, LOC showed a non-significant trend towards *anti-adaptation*; that is, increased (rather than decreased) response for category repetitions [$t(14) = 1.7, \text{two-tailed } p = 0.11$]. This may reflect the deployment of additional attention towards the objects within a scene when category is repeated.

The failure to observe a significant category-related fMRIa effect in any region except the left medial PPA is striking, especially given that we can decode outdoor categories with high accuracy in all of our ROIs. In contrast, landmark repetition effects in the PPA and RSC were robust. Indeed, direct comparison revealed that landmark adaptation was stronger than category adaptation in PPA [$t(14) = 1.73, p = 0.05$] and RSC [$t(14) = 2.7, p = 0.009$]. Even in the left medial PPA region that showed the strongest category adaptation effect, the landmark adaptation effect was numerically greater, although the difference was not significant [$t = 1.1, n.s.$]. This contrasts sharply with the MVPA findings, which suggested that PPA neural codes are equally informative about Penn landmarks and outdoor categories.

3.5. Spatial distribution of effects within ROIs

The previous results suggest a clear disjunction between the neural mechanisms that contribute to MVPA and the neural

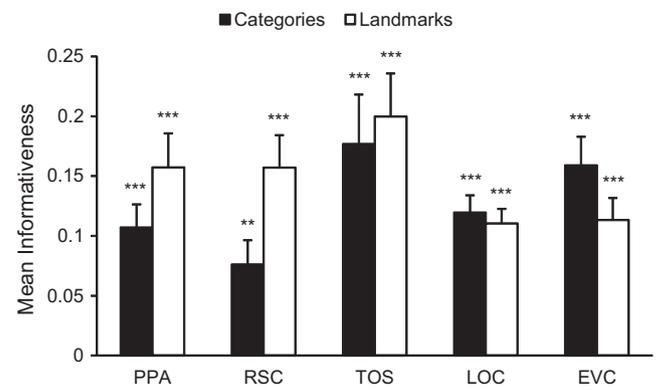


Fig. 7. Average voxelwise informativeness (mean \pm SEM) for each ROI. Informativeness was defined as the cross-run correlation between response levels for all 10 categories or all 10 landmarks. Consistent with the MVPA classification results, mean informativeness was above chance in all regions. Informativeness was greater for landmarks than for categories in RSC but did not differ between categories and landmarks in any other region. ** $p < 0.01$; *** $p < 0.001$.

mechanisms that contribute to fMRIa. In order to better understand the relationship between these two mechanisms, we tested whether the voxels that showed adaptation were the same voxels that contributed to MVPA decoding.

To answer this question, we needed to quantify the informativeness of the activation levels for each individual voxel. We adopted a measure developed by previous researchers (Kravitz et al., 2010; Mitchell et al., 2008): the between-run correlation of response values for each voxel. The logic of this measure is straightforward: if the response values of a voxel convey information about category (or landmark) then these response values should be reliable across runs, and between-run correlation should be high. On the other hand, if the response values are merely noise, then they should be unreliable across runs, and between-run correlation should be low. Note that this reasoning mimics the logic of the pattern classification scheme used for our MVPA analysis, but with one important difference: whereas in MVPA we assess the reliability of response levels across many voxels for a given stimulus category (or landmark), here we assess the reliability of response levels across many stimulus categories (or landmarks) for a given voxel.

To validate this approach, we calculated informativeness values averaged across all voxels in our various ROIs (Fig. 7) Average informativeness was above chance in all regions for both outdoor categories [PPA $t(14) = 5.5$; RSC $t(14) = 3.7$, TOS $t(12) = 4.3$, LOC $t(14) = 8.2$, EVC $t(13) = 6.6$, all $ps < 0.002$] and Penn landmarks [PPA $t(14) = 5.5$; RSC $t(14) = 5.8$, TOS $t(12) = 5.5$, LOC $t(14) = 9.1$, EVC $t(13) = 6.2$, all $ps < 0.0005$], consistent with previous findings that both landmarks and categories can be decoded with a high degree of accuracy. Informativeness values for Penn landmarks vs. Outdoor Categories roughly tracked classification performance in the PPA and RSC. Specifically, there was no significant difference between landmark and category informativeness in the PPA [$t(14) = 1.4, p = 0.17$ two-tailed], while informativeness values were higher for landmarks than for categories in RSC [$t(14) = 2.5, p = 0.02$ two-tailed].

We next examined whether the voxels that were highly informative about landmark identity or scene category were same the voxels that showed reduced response when these quantities were repeated. To do this, we examined the correlation between the informativeness values and adaptation values across all voxels within each ROI. We observed a significant correlation between landmark informativeness and landmark adaptation in the PPA [mean $r = -0.10, t(14) = -2.7, p = 0.009$] and RSC [mean $r = -0.11, t(14) = -2.1, p = 0.03$]. In contrast, there was no significant correlation between category informativeness and category adaptation in

either of these ROIs [$t_s < 1$, n.s.]. These null results probably reflect the fact that category adaptation effects were not significant in these regions. When we examined the left medial PPA region that show significant category adaptation, there was a significant correlation between category adaptation and category informativeness [mean $r = -0.13$, $t(14) = -2.0$, $p = 0.04$].

These results suggest that voxels that convey information about either landmark identity or scene category in their response levels also exhibit adaptation when these quantities are repeated. The mechanisms that support voxelwise encoding and fMRIa appear to be, at least to some extent, physically coterminous in the PPA and RSC.

3.6. Whole-brain analyses of MVPA classification and fMRI adaptation

To determine whether any region outside of predefined ROIs exhibits above-chance classification for outdoor categories or Penn landmarks, we performed a “searchlight” analysis, which allowed us to examine classification performance in the neighborhood surrounding each voxel of the brain. Results are shown in Fig. 8a. As can be seen, classification performance was quite high for both stimulus classes throughout many regions of the occipital, temporal, and parietal lobes. Beyond the functional ROIs defined earlier (compare Fig. 8a–c), we also observed high classification performance in ventral stream regions posterior to the PPA and LOC, and high classification performance for landmarks in the intraparietal sulcus (superior to TOS). Classification in ventral stream regions might reflect processing of intermediate-level visual features such as color, while in parietal regions may reflect processing of the spatial aspects of the stimuli. Although less prominent, there are also small patches of high classification performance in the frontal lobes, which could reflect semantic or verbal recoding of the stimuli. Overall, it is notable that classification performance was high in a wide range of visually-responsive regions, a finding that likely reflects the fact that there are many different feature dimensions along which scene categories and individual landmarks can be reliably distinguished.

We also performed a whole-brain analysis of the fMRI adaptation effects. Landmark repetition led to reduced response (adaptation) in a smaller set of regions, including the PPA and RSC, and adjoining territory in the lingual gyrus and retrosplenial cortex proper (Fig. 8b; see Morgan et al., 2011, for additional details). Thus, the set of regions showing adaptation for landmarks differs substantially from the set of regions showing high MVPA classification performance. Most notably, adaptation was strongest in medial parietal regions, and was much weaker or nonexistent in posterior visual regions showing the highest classification performance. The PPA and RSC are an area of overlap, within which both adaptation and classification are significant.

No significant category-related adaptation effects were observed at the $p < 0.01$ uncorrected threshold in any region (data not shown). The failure to observe category-related adaptation in any brain region may seem surprising in light of previous studies reporting response reductions in the fusiform gyrus and frontal lobe regions when different exemplars of the same object category are repeated. However, it is worth noting that these studies utilized a “neural priming” paradigm in which items were repeated over longer intervals with several intervening items. We have previously speculated that such “long-interval” repetition regimes might induce neural adaptation mechanisms than are fundamentally different than those induced by the “medium-interval” repetitions examined here (Epstein, Parker, & Feiler, 2008). We take up the issue of different fMRI adaptation mechanisms further in the Discussion.

4. Discussion

The current study used MVPA and fMRIa to examine the neural codes that support recognition of visual scenes. We addressed two main issues. First, to what extent do the PPA and RSC support recognition of scenes at either the categorical or the individual exemplar level? Second, to what extent are the representational distinctions revealed by MVPA consistent with the representational distinctions revealed by fMRIa? Our data suggest that the first question cannot be fully answered without also addressing the second. In the discussion below, we will first discuss the MVPA data on category vs. exemplar encoding, and then discuss how the fMRIa data shades our interpretation of the MVPA results.

4.1. MVPA findings on category vs. landmark encoding

When looking at a visual scene, such as an image of a kitchen or a beach, one can either identify it at the categorical level (“kitchen”, “beach”) or at the exemplar level (“the kitchen of the Penn Center for Cognitive Neuroscience”, “Vanderbilt Beach in Naples Florida”). As these descriptions indicate, scenes defined categorically have no specific locations in the world, while scenes defined as specific exemplars have the potential to be associated with specific spatial coordinates. Thus, the issue of representational level relates intimately to the putative function of scene-responsive regions. Categorical recognition is likely to be more useful for understanding the kind of objects and actions that should be expected within the environment, while exemplar recognition is likely to be more useful for identifying a scene as a specific location during spatial navigation.

We addressed this issue by examining multi-voxel patterns associated either with general scene categories (beach, jungle, etc.) or specific scene exemplars drawn from the Penn campus. (Note that in this usage, “exemplar” refers to a specific place or location in the world, rather than to a specific image.) Our results indicated that both categorical and exemplar information could be decoded at rates well above chance in both the PPA and RSC, as well as in several other cortical regions. To our knowledge, this is the first study to directly compare MVPA performance across these two distinct levels of representation. Although there are some suboptimalities to our design—most notably, the fact that Penn landmarks were personally familiar to the subjects while the locations depicted in the outdoor category images were not, and the fact that Penn landmarks and outdoor categories were not shown in the same scan runs—these results do provide some evidence that PPA and (to a lesser extent) RSC might be involved in both levels of scene recognition (although see Section 4.2 below).

Our data also revealed some intriguing differences between the PPA and RSC. In the PPA, there was little evidence that one stimulus class was favored over the other: MVPA classification performance, average Euclidean distance between MVPs, and average informativeness of individual PPA voxels was equivalent for outdoor categories and Penn landmarks. RSC, on the other hand, showed a preference for the Penn landmark stimuli: classification performance was better for the landmarks than for the outdoor categories, Euclidean distances between different landmarks were larger than Euclidean distances between different categories, and mean voxelwise informativeness was higher for landmarks than for categories. These findings are consistent with previous reports that PPA is more involved in the visual recognition of scenes while RSC is more involved in calculating spatial quantities associated with the locations depicted in scenes (Epstein, 2008; Epstein, Parker, et al., 2007; Park & Chun, 2009). These spatial quantities would be more salient and varied for the Penn landmarks than for the outdoor categories, and thus we would expect RSC to consider the Penn

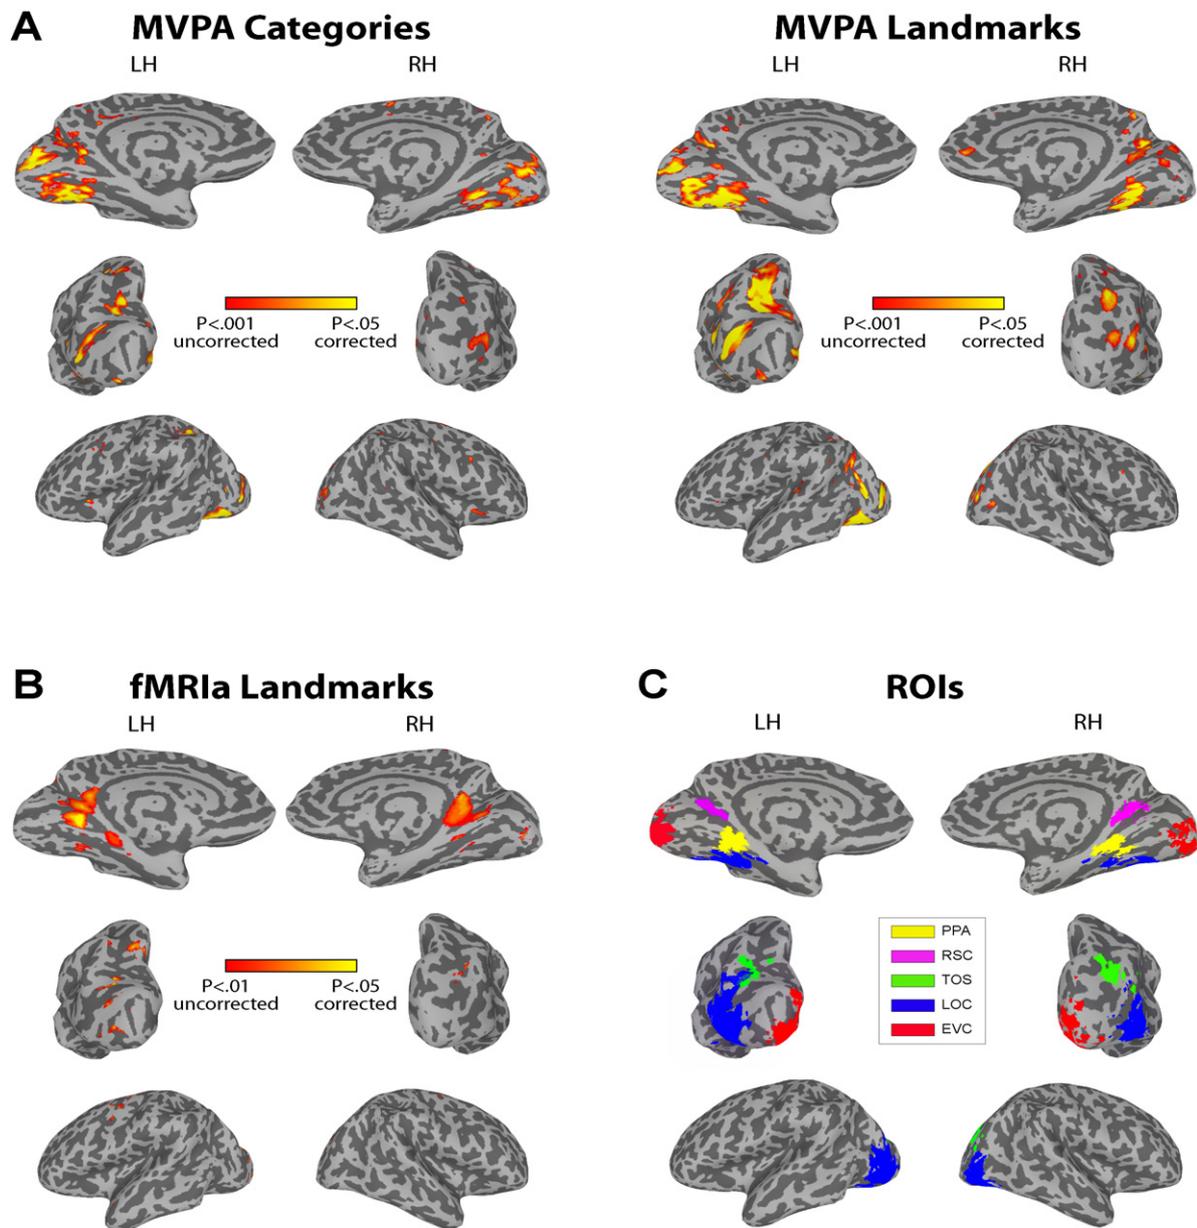


Fig. 8. Whole-brain analyses. (A) MVPA searchlight analysis revealed a wide swath of territory in occipito-temporal-parietal cortex for which multi-voxel activity patterns conveyed information about scene category (left) or landmark identity (right). Orange voxels are significant at $p < 0.001$ uncorrected; yellow voxels are significant at $p < 0.05$ corrected for multiple comparisons (a more stringent threshold). Note that the medial views are tilted slightly to expose the ventral side. (B) fMRI adaptation effects induced by landmark repetitions were generally confined to scene-responsive ROIs and nearby territory. Adaptation effects induced by category repetition were not significant in any area of the brain at these thresholds (not shown). (C) Functional ROIs. Boundaries reflect the across-subject ROI intersection that most closely matches the average size of the individual subject ROIs.

landmarks to be the more representationally disparate stimulus class.

The finding that PPA considered our 10 outdoor categories to be no more representationally distinct than our 10 Penn landmarks is potentially a puzzling one. Previous accounts have suggested that the PPA might represent visual (Cant & Goodale, 2007), geometric (Epstein & Kanwisher, 1998; Park et al., 2011), or semantic (Bar & Aminoff, 2003) aspects of scenes. The outdoor categories are more visually and semantically disparate than the Penn landmarks, so one might expect that the representational gamut would be larger for the outdoor categories. But this was not what we observed—average Euclidean distances between patterns did not differ between the Penn landmarks and outdoor categories in PPA. Nor did we see a relationship between visual and neural similarity. Although the visual features examined in this analysis are

admittedly quite low-level—and did not include color information, which might be important for scene recognition—the absence of neural–visual relationship is still somewhat surprising, given that computational work suggests that at least some high-level scene properties are correlated with low-level visual statistics (Torralba & Oliva, 2003).

These MVPA data tend to support a variant of a “categorical” view of scene representation in the PPA under which different landmarks are considered, on average, to be as representationally-distinct as different categories. We speculate that this finding might depend, in part, on the fact that subjects were highly familiar with the Penn landmarks. Previous behavioral work on object recognition suggest that highly familiar items tend to be identified at the individual exemplar level while unfamiliar items tend to be recognized at the basic categorical level (Tanaka &

Taylor, 1991). Analogously, we hypothesize that once a landmark or scene becomes familiar, the PPA might treat it as a distinct “category” for purposes of recognition. Previous neuroimaging studies have demonstrated that navigational experience can affect PPA response: the PPA responds more strongly to familiar vs. unfamiliar places (Epstein, Higgins, et al., 2007) and more strongly to navigationally-relevant vs. non-navigationally relevant objects (Janzen & van Turenout, 2004). It is reasonable to suppose that familiarity might modify not just the level of response in the PPA but also the structure of the underlying representational code. Under this account, one might expect to see a more hierarchical representational organization for unfamiliar exemplars, with narrowly-tuned exemplars encompassed by wider categories – a point that should be explored in future experiments. In addition, it would be worthwhile to examine MVPs for categories and exemplars interspersed within the same runs, as it is possible that the landmark vs. category equivalence observed in our gamut analysis reflects dynamic remapping of the gamut for each run rather than a true equivalence in representational space (Panis, Wagemans, & Op de Beeck, 2011).

Our data did not reveal an organizational principle behind PPA coding of outdoor categories and familiar landmarks. Despite this, we suspect that such a principle must exist. Inspection of the confusion matrices (Fig. 5a) reveals that there is considerable structure in the off-diagonal elements. Although we cannot assess whether this off-diagonal structure is reliable, its presence suggests that the PPA considers some scene categories and landmarks to be more similar than others, rather than considering all such items to be equally distinct. We can only speculate about the nature of the underlying similarity metric, which did not seem to correspond to similarities in low-level features. Previous work suggests that PPA response is strongly affected by geometric quantities such as openness or closedness (Park et al., 2011) or the principal axis of the scene (Epstein, 2008; Shelton & Pippitt, 2007) and previous MVPA studies have shown that PPA response patterns cluster by geometric similarity (Kravitz et al., 2010; Park et al., 2011). We suggest that classification performance in the current experiment might be driven in part by differences in the geometric features of scenes, which might vary equivalently for the outdoor categories and the Penn landmarks. Alternatively, we cannot exclude the possibility that the PPA encodes a semantic space, in which different landmarks and different categories are related to each other based on non-physical features.

We also observed above-chance MVPA classification performance in posterior visual regions (EVC) and object-selective cortex (LOC). These findings are consistent with results of previous studies (Walther et al., 2009) and are not surprising given the existence of reliable visual differences between the landmarks and outdoor categories. Notably, our model of visual dissimilarity predicted a significant fraction of the neural dissimilarity between patterns in both EVC and LOC, a relationship that was not found in PPA or RSC. We speculate that EVC might encode simple visual features that differ reliably between scenes, while LOC might encode characteristic objects or object-based features that are also predictive of scene category and identity (MacEvoy & Epstein, 2011).

4.2. Relating MVPA findings to fMRI data

Whereas MVPA indicated that scenes can be decoded in the PPA, RSC, TOS, LOC and EVC, fMRI results suggested that scene information was restricted to a much smaller set of cortical regions. In particular, landmark adaptation was observed in “scene” regions (PPA, RSC, and TOS) but not “object” regions (LOC) or early visual cortex (EVC). Even more strikingly, category adaptation was only observed in the left medial subregion of the PPA, with no hint of a repetition suppression effect in any other area.

What are we to make of the apparent inconsistencies between the MVPA and the fMRI data? Although one could argue that fMRI is simply less sensitive to representational distinctions than MVPA (Sapountzis, Schluppeck, Bowtell, & Peirce, 2010), this cannot explain the data from the PPA: here MVPA found Penn landmarks and outdoor categories to be equally decodable, whereas fMRI found a stronger effect for landmark repetition than for category repetition. Rather, we believe that these results are consistent with earlier findings suggesting that fMRI and MVPA might interrogate different aspects of the neural code (Drucker & Aguirre, 2009). We propose three hypotheses about the underlying mechanisms that may drive these two effects (see Fig. 9).

The first hypothesis, adopted directly from Drucker and Aguirre (2009), is that fMRI reflects the tuning of individual neurons (or perhaps, individual cortical columns) while MVPA reflects clustering at a coarser anatomical scale. In this view, PPA neurons would be tuned to specific landmarks or scenic exemplars, but these neurons would be clustered according to categorical or geometric similarity, thus permitting decoding of both landmarks and categories using multivoxel patterns. The much weaker adaptation effect for category might reflect an absence of categorically-tuned neurons, except perhaps in the left medial PPA. Similarly, RSC might contain neurons tuned for individual landmarks but not for categories, thus leading to adaptation only for the landmarks, while LOC and EVC neurons might be tuned for simpler features that are not consistent across different exemplars of a scenic category or different views of a scene exemplar, thus leading to an absence of adaptation for both stimulus classes.

This interpretation of the fMRI results in terms of neural (or columnar) tuning runs counter to our previous interpretation of such results in terms of adaptation at the synaptic inputs to a neuron (Epstein et al., 2008). One important difference between the current design and previous experiments on scene adaptation is the length of the repetition interval, which was 100–700 ms in our previous experiments, compared to 2 s here. It is possible that “short-interval” (100–700 ms) and “medium interval” (2–3 s) repetitions might elicit adaptation through different mechanisms. Previous studies using short-interval repetition have found adaptation effects that are viewpoint- and stimulus-specific (Epstein et al., 2003, 2005; Fang, Murray, & He, 2007); in contrast, here we observed some degree of viewpoint-tolerance (and even some degree of generalization across category exemplars in the left medial PPA). Although this viewpoint-tolerance might be explained simply by the high degree of overlap between the images corresponding to each Penn landmark, it is also possible that it reflects the workings of an adaptation mechanism that operates at a later processing stage, such as the unit or column, rather than inputs to a unit or column. Neurophysiological evidence suggests that short-interval adaptation operates on synaptic inputs, as evidenced by adaptation effects that are more stimulus-specific than the neuronal response (De Baene & Vogels, 2010; Sawamura, Orban, & Vogels, 2006). To our knowledge, this hyperspecificity of adaption has not been tested for medium-interval repetition.

The second hypothesis is that fMRI reflects adaptation at the synaptic inputs even for the medium-interval repetitions used in this experiment, while MVPA reflects neuronal outputs. Under this account, fMRI would be greater for Penn landmarks than for outdoor categories, because different views of the same landmark activate partially overlapping inputs, while different exemplars of the same scene category do not. In addition to the neurophysiological data outlined above, this hypothesis is further supported by a recent study of adaptation effects in monkey IT, which found that response reduction was only observed in the first 300 ms of the response but not in the later components (Liu, Murray, & Jagadeesh, 2009). Although we must be careful when generalizing from monkeys to humans, and from object-selective to

MVPA vs. Adaptation: Three Hypotheses

Mechanisms underlying MVPA in *red*; mechanisms underlying adaptation in *blue*.

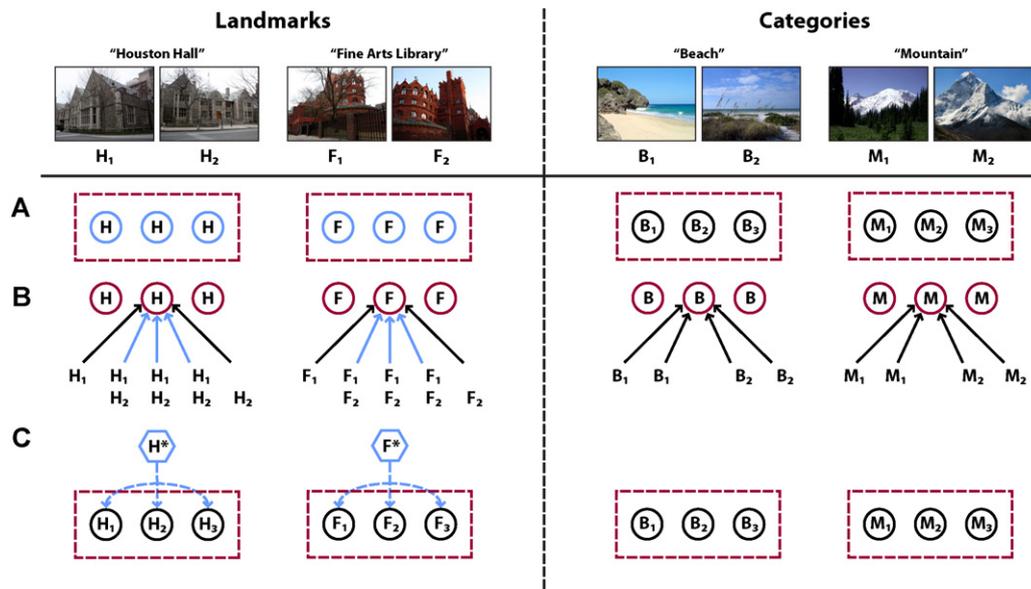


Fig. 9. Three hypotheses about the neural mechanisms that underlie MVPA and fMRI adaptation in the PPA. Units (which might be either neurons or columns) are represented by circles; synaptic inputs to units are represented by solid arrows; dashed boxes represent coarse-scale groupings of units; dashed arrows represent transient coalitions between units. Elements that drive MVPA are in red; elements that drive fMRI adaptation are in blue; elements that drive neither are in black. (A) Under the first hypothesis, adaptation operates on individual units (blue circles) and thus reflects neuronal (or columnar) tuning, while MVPA reflects coarse-scale groupings of units (red dashed boxes). Adaptation is observed for landmarks but not categories because neurons are selective for individual landmarks (H, F) and individual scene exemplars (B_1 , B_2) but not for scene categories (B, M). (B) Under the second hypothesis, adaptation operates on the inputs to each unit (blue arrows), while MVPA reflects neuronal tuning (red circles). Adaptation is observed for landmarks but not categories because different views of the same landmark (H_1 , H_2) activate overlapping inputs while different exemplars of the same category (B_1 , B_2) do not. (C) Under the third hypothesis, adaptation reflects the formation of a transient coalition of units (blue dashed lines), possibly coordinated by top-down inputs from other regions (blue hexagon), while MVPA reflects a more enduring, coarse-scale topographical organization (red dashed boxes). Adaptation is observed for landmarks but not categories because only landmark repetitions are fulfilling of expectations and thus lead to quicker re-instantiation of a neural coalition. In this scenario, neither fMRI nor MVPA directly index neuronal tuning.

scene-selective regions (Weiner, Sayres, Vinberg, & Grill-Spector, 2010), these data are consistent with the idea that fMRI interrogates the inputs and initial response to a stimulus rather than the ultimate outputs. In the case of the PPA, one might suppose that view-specific inputs are converted to a more “abstract” code, in which different scene categories and different landmarks are represented independent of their visual qualities.

The third hypothesis is that MVPA reveals coarse-grain clustering of features, while fMRI reflects dynamic processes that operate on top of the underlying neural code. For example, adaptation might reflect the facility with which the system creates transient neuronal coalitions that link together the features that correspond to a given landmark or category. These coalitions might be local to the PPA and RSC, or they might involve interaction between these regions and higher-level areas in the frontal lobe, hippocampus, or retrosplenial cortex proper (BA 29/30). This hypothesis builds on theoretical work suggesting that visual recognition involves an interplay between bottom-up input and top-down interpretation (Friston, 2005), a view that gains support from a recent finding that fMRI effects are larger when repetitions are more frequent and thus more fulfilling of perceptual expectations (Summerfield, Trittschuh, Monti, Mesulam, & Egner, 2008) (but see Kaliukhovich & Vogels, 2010). It is reasonable to suppose that “expectation” in the current experiment would work on the level of scene exemplars rather than scene categories. That is, viewing a given scene or landmark leads one to expect that one will encounter visual features corresponding to that scene or landmark in the immediate future. Because different images of the same landmark share more visual features than different images of the same scene category, landmark repetitions might have been treated as more fulfilling of

expectations than category repetitions. The end results would be stronger fMRI for the Penn Landmarks than for the outdoor categories. Note that whereas the second hypothesis proposes that adaptation occurs *early* in the neuronal response to a stimulus, this hypothesis proposes that adaptation occurs *late*.

These three accounts make different predictions which could be potentially tested in further fMRI experiments. In particular, the first and third accounts posit that the representations revealed by fMRI are more directly tied to recognition than the representations revealed by MVPA – either because fMRI indexes neuronal tuning directly, or because it indexes dynamic processes that are the mechanism by which recognition operates. In contrast, the second account posits that the representations revealed by MVPA should be more closely tied to recognition, because these reflect neuronal outputs rather than synaptic inputs. Thus, one way to adjudicate between the three accounts would be to examine whether the representational distinctions revealed by fMRI or MVPA more closely relate to the representational distinctions revealed by behavior. Another issue of potential importance is the timecourse of the fMRI effect: the second account suggests that it operates on the early component of the neuronal response, while the third account suggests that it operates on the late components. These predictions could be tested by varying the length of the stimulus presentation, and also by using pattern masks to selectively interrupt later, top-down response components. Finally, the second account proposes that MVPA reflects neuronal or columnar tuning while the first and third account propose that it reflects organization at a coarser spatial scale. Several authors have proposed methods for addressing this issue (Freeman, Brouwer, Heeger, & Merriam, 2011; Kamitani & Tong, 2005; Sasaki et al., 2006; Swisher et al., 2010) – for

example, by examining whether classification performance is reduced by spatial smoothing (Op de Beeck, 2010). If classification performance is unaffected by spatial smoothing, this would argue in favor of the first or third account, under which MVPA reflects coarse-scale organization. On the other hand, if spatial smoothing reduces classification performance, then the second account, under which MVPA directly indexes neuronal tuning, become more plausible.

Finally, we note that the spatial distribution of fMRI and MVPA effects across brain regions might provide information that could partially adjudicate the three accounts. Inspection of Fig. 8 suggests that the brain regions that exhibit high classification performance tend to be “earlier” along the visual processing stream than regions that exhibit adaptation. The pattern would be consistent with scenario 1, because earlier visual regions would be expected to represent category- and landmark-distinguishing visual features in a spatially-coarse manner that would be easily read out by MVPA. In contrast, higher-level regions, which would be more likely to explicitly encode category and landmark identity at the neuronal level, would likely support spatially interdigitated representations of these quantities that are harder to decode from multi-voxel activity patterns. This pattern is also consistent with scenario 3, in which adaptation operates through top-down signals, and thus would likely be more evident in “higher-level” than in “low-level” processing regions.

5. Conclusion

We used MVPA and fMRI to investigate the neural codes that underlie scene recognition. We were especially interested in identifying neural codes corresponding to the coding of scene categories and individual scene exemplars (in this case, individual landmarks from the Penn campus). Data from both MVPA and fMRI are in agreement that PPA and RSC represent scenes at the exemplar level. However, these two analysis techniques gave inconsistent results when it comes to the coding of scene categories: whereas MVPA strongly suggests that PPA and RSC encode category information, fMRI suggests that PPA only encodes category information in the left medial subregion and RSC does not encode category information at all. These data suggest that MVPA and fMRI interrogate different aspects of the neuronal response. Given that these techniques are used frequently to make claims about representations supported by different brain regions, and indeed have become part of the central toolkit of cognitive neuroscience, we believe that it is critical to more precisely delineate the neuronal signals that underlie these two techniques.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.neuropsychologia.2011.09.042.

References

- Aguirre, G. K. (2007). Continuous carry-over designs for fMRI. *Neuroimage*, 35(4), 1480–1494.
- Arcaro, M. J., McMains, S. A., Singer, B. D. & Kastner, S. (2009). Retinotopic organization of human ventral visual cortex. *Journal of Neuroscience*, 29(34), 10638–10652.
- Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, 5(8), 617–629.
- Bar, M. & Aminoff, E. (2003). Cortical analysis of visual context. *Neuron*, 38, 347–358.
- Biederman, I. (1972). Perceiving real-world scenes. *Science*, 177(43), 77–80.
- Cant, J. S. & Goodale, M. A. (2007). Attention to form or surface properties modulates different regions of human occipitotemporal cortex. *Cerebral Cortex*, 17(3), 713–731.
- Cox, D. D. & Savoy, R. L. (2003). Functional magnetic resonance imaging (fMRI) “brain reading”: Detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage*, 19(2 Part 1), 261–270.
- De Baene, W. & Vogels, R. (2010). Effects of adaptation on the stimulus selectivity of macaque inferior temporal spiking activity and local field potentials. *Cerebral Cortex*, 20(9), 2145–2165.
- Drucker, D. M. & Aguirre, G. K. (2009). Different spatial scales of shape similarity representation in lateral and ventral LOC. *Cerebral Cortex*, 19(10), 2269–2280.
- Epstein, R., DeYoe, E. A., Press, D. Z., Rosen, A. C. & Kanwisher, N. (2001). Neuropsychological evidence for a topographical learning mechanism in parahippocampal cortex. *Cognitive Neuropsychology*, 18(6), 481–508.
- Epstein, R., Graham, K. S. & Downing, P. E. (2003). Viewpoint-specific scene representations in human parahippocampal cortex. *Neuron*, 37, 865–876.
- Epstein, R. & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, 392(6676), 598–601.
- Epstein, R. A. (2008). Parahippocampal and retrosplenial contributions to human spatial navigation. *Trends in Cognitive Sciences*, 12(10), 388–396.
- Epstein, R. A. & Higgins, J. S. (2007). Differential parahippocampal and retrosplenial involvement in three types of visual scene recognition. *Cerebral Cortex*, 17(7), 1680–1693.
- Epstein, R. A., Higgins, J. S., Jablonski, K. & Feiler, A. M. (2007). Visual scene processing in familiar and unfamiliar environments. *Journal of Neurophysiology*, 97(5), 3670–3683.
- Epstein, R. A., Higgins, J. S. & Thompson-Schill, S. L. (2005). Learning places from views: Variation in scene processing as a function of experience and navigational ability. *Journal of Cognitive Neuroscience*, 17(1), 73–83.
- Epstein, R. A., Parker, W. E. & Feiler, A. M. (2007). Where am I now? Distinct roles for parahippocampal and retrosplenial cortices in place recognition. *Journal of Neuroscience*, 27(23), 6141–6149.
- Epstein, R. A., Parker, W. E. & Feiler, A. M. (2008). Two kinds of fMRI repetition suppression? Evidence for dissociable neural mechanisms. *Journal of Neurophysiology*, 99(6), 2877–2886.
- Fang, F., Murray, S. O. & He, S. (2007). Duration-dependent fMRI adaptation and distributed viewer-centered face representation in human visual cortex. *Cerebral Cortex*, 17(6), 1402–1411.
- Fei-Fei, L., Iyer, A., Koch, C. & Perona, P. (2007). What do we perceive in a glance of a real-world scene? *Journal of Vision*, 7(1), 10.
- Freeman, J., Brouwer, G. J., Heeger, D. J. & Merriam, E. P. (2011). Orientation decoding depends on maps, not columns. *Journal of Neuroscience*, 31(13), 4792–4804.
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, 360(1456), 815–836.
- Greene, M. R. & Oliva, A. (2009). Recognition of natural scenes from global properties: Seeing the forest without representing the trees. *Cognitive Psychology*, 58(2), 137–176.
- Grill-Spector, K., Henson, R. & Martin, A. (2006). Repetition and the brain: Neural models of stimulus-specific effects. *Trends in Cognitive Sciences*, 10(1), 14–23.
- Grill-Spector, K. & Malach, R. (2001). fMR-adaptation: A tool for studying the functional properties of human cortical neurons. *Acta Psychologica*, 107(1–3), 293–321.
- Habib, M. & Sirigu, A. (1987). Pure topographical disorientation - A definition and anatomical basis. *Cortex*, 23(1), 73–85.
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L. & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, 293(5539), 2425–2430.
- Hung, C. P., Kreiman, G., Poggio, T. & DiCarlo, J. J. (2005). Fast readout of object identity from macaque inferior temporal cortex. *Science*, 307, 863–866.
- Janzen, G. & van Turenhout, M. (2004). Selective neural representation of objects relevant for navigation. *Nature Neuroscience*, 7(6), 673–677.
- Kaliukhovich, D. A. & Vogels, R. (2010). Stimulus repetition probability does not affect repetition suppression in macaque inferior temporal cortex. *Cerebral Cortex*, 20(10), 2425–2430.
- Kamitani, Y. & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, 8(5), 679–685.
- Kourtzi, Z. & Kanwisher, N. (2001). Representation of perceived object shape by the human lateral occipital complex. *Science*, 293(5534), 1506–1509.
- Kravitz, D., Peng, C. & Baker, C. I. (2010). The structure of scene representations across the ventral visual pathway. *Journal of Vision*, 10(7), 1224.
- Kriegeskorte, N., Goebel, R. & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences of the United States of America*, 103(10), 3863–3868.
- Liu, Y., Murray, S. O. & Jagadeesh, B. (2009). Time course and stimulus dependence of repetition-induced response suppression in inferotemporal cortex. *Journal of Neurophysiology*, 101(1), 418–436.
- MacEvoy, S. P. & Epstein, R. A. (2011). Constructing scenes from objects in human occipitotemporal cortex. *Nature Neuroscience*, 14(10), 1323–1329.
- Mendez, M. F. & Chierri, M. M. (2003). Agnosia for scenes in topographagnosia. *Neuropsychologia*, 41(10), 1387–1395.
- Mitchell, T. M., Shinkareva, S. V., Carlson, A., Chang, K. M., Malave, V. L., Mason, R. A., et al. (2008). Predicting human brain activity associated with the meanings of nouns. *Science*, 320(5880), 1191–1195.
- Morgan, L. K., Macevoy, S. P., Aguirre, G. K. & Epstein, R. A. (2011). Distances between real-world locations are represented in the human hippocampus. *Journal of Neuroscience*, 31(4), 1238–1245.
- Nichols, T. E. & Holmes, A. P. (2002). Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Human Brain Mapping*, 15(1), 1–25.
- Norman, K. A., Polyn, S. M., Detre, G. J. & Haxby, J. V. (2006). Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*, 10(9), 424–430.
- Op de Beeck, H. P. (2010). Probing the mysterious underpinnings of multi-voxel fMRI analyses. *Neuroimage*, 50(2), 567–571.

- Panis, S., Wagemans, J. & Op de Beeck, H. P. (2011). Dynamic norm-based encoding for unfamiliar shapes in human visual cortex. *Journal of Cognitive Neuroscience*, 23(7), 1829–1843.
- Park, S., Brady, T. F., Greene, M. R. & Oliva, A. (2011). Disentangling scene content from spatial boundary: Complementary roles for the parahippocampal place area and lateral occipital complex in representing real-world scenes. *Journal of Neuroscience*, 31(4), 1333–1340.
- Park, S. & Chun, M. M. (2009). Different roles of the parahippocampal place area (PPA) and retrosplenial cortex (RSC) in panoramic scene perception. *Neuroimage*, 47(4), 1747–1756.
- Potter, M. C. (1975). Meaning in visual search. *Science*, 187(4180), 965–966.
- Renninger, L. W. & Malik, J. (2004). When is scene identification just texture recognition? *Vision Research*, 44, 2301–2311.
- Sapountzis, P., Schluppeck, D., Bowtell, R. & Peirce, J. W. (2010). A comparison of fMRI adaptation and multivariate pattern classification analysis in visual cortex. *Neuroimage*, 49(2), 1632–1640.
- Sasaki, Y., Rajimehr, R., Kim, B. W., Ekstrom, L. B., Vanduffel, W. & Tootell, R. B. (2006). The radial bias: A different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron*, 51(5), 661–670.
- Sawamura, H., Orban, G. A. & Vogels, R. (2006). Selectivity of neuronal adaptation does not match response selectivity: A single-cell study of the fMRI adaptation paradigm. *Neuron*, 49(2), 307–318.
- Sewards, T. V. (2010). Neural structures and mechanisms involved in scene recognition: A review and interpretation. *Neuropsychologia*, 49(3), 277–298.
- Shelton, A. L. & Pippitt, H. A. (2007). Fixed versus dynamic orientations in environmental learning from ground-level and aerial perspectives. *Psychological Research*, 71(3), 333–346.
- Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M. M. & Eger, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*.
- Swisher, J. D., Gatenby, J. C., Gore, J. C., Wolfe, B. A., Moon, C. H., Kim, S. G., et al. (2010). Multiscale pattern analysis of orientation-selective activity in the primary visual cortex. *Journal of Neuroscience*, 30(1), 325–330.
- Takahashi, N., Kawamura, M., Shiota, J., Kasahata, N. & Hirayama, K. (1997). Pure topographic disorientation due to right retrosplenial lesion. *Neurology*, 49(2), 464–469.
- Tanaka, J. W. & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder. *Cognitive Psychology*, 23(3), 457–482.
- Torralba, A. & Oliva, A. (2003). Statistics of natural image categories. *Network*, 14(3), 391–412.
- Tversky, B. & Hemenway, K. (1983). Categories of environmental scenes. *Cognitive Psychology*, 15, 121–149.
- Walther, D. B., Caddigan, E., Fei-Fei, L. & Beck, D. M. (2009). Natural scene categories revealed in distributed patterns of activity in the human brain. *Journal of Neuroscience*, 29(34), 10573–10581.
- Weiner, K. S., Sayres, R., Vinberg, J. & Grill-Spector, K. (2010). fMRI-adaptation and category selectivity in human ventral temporal cortex: Regional differences across time scales. *Journal of Neurophysiology*, 103(6), 3349–3365.
- Williams, M. A., Dang, S. & Kanwisher, N. G. (2007). Only some spatial patterns of fMRI response are read out in task performance. *Nature Neuroscience*, 10(6), 685–686.

Farm Exemplars



Huntsman Hall Exemplars

