

Distributed Category-Specific Recognition-Memory Signals in Human Perirhinal Cortex

Chris B. Martin,^{1,2} Rosemary A. Cowell,³ Paul L. Gribble,^{1,4}
 Jessey Wright,^{5,6} and Stefan Köhler^{1,7*}

ABSTRACT: Evidence from a large body of research suggests that perirhinal cortex (PrC), which interfaces the medial temporal lobe with the ventral visual pathway for object identification, plays a critical role in item-based recognition memory. The precise manner in which PrC codes for the prior occurrence of objects, however, remains poorly understood. In the present functional magnetic resonance imaging (fMRI) study, we used multivoxel pattern analyses to examine whether the prior occurrence of faces is coded by distributed patterns of PrC activity that consist of voxels with decreases as well as increases in signal. We also investigated whether pertinent voxels are preferentially tuned to the specific object category to which judged stimuli belong. We found that, when no a priori constraints were imposed on the direction of signal change, activity patterns that allowed for successful classification of recognition-memory decisions included some voxels with decreases and others with increases in signal in association with perceived prior occurrence. Moreover, successful classification was obtained in the absence of a mean difference in activity across the set of voxels in these patterns. Critically, we observed a positive relationship between classifier accuracy and behavioral performance across participants. Additional analyses revealed that voxels carrying diagnostic information for classification of memory decisions showed category specificity in their tuning for faces when probed with an independent functional localizer in a nonmnemonic task context. These voxels were spatially distributed in PrC, and extended beyond the contiguous voxel clusters previously described as the anterior temporal face patch. Our findings provide support for proposals, recently raised in the neurophysiological literature, that the prior occurrence of objects is coded by distributed PrC representations. They also suggest that the stimulus category to which an item belongs shapes the organization of these distributed representations.
 © 2015 Wiley Periodicals, Inc.

KEY WORDS: medial temporal lobe; familiarity; faces; fMRI; multivoxel pattern analysis

INTRODUCTION

A large body of research findings in humans, nonhuman primates, and rodents converge to suggest that perirhinal cortex (PrC), a structure in the medial temporal lobe (MTL) that is intimately connected with the ventral visual pathway for object identification, plays a critical role in the ability to discriminate between previously encountered and novel stimuli (i.e., recognition of prior occurrence; Meunier et al., 1993; Eacott et al., 1994; Brown and Aggleton, 2001; Eichenbaum et al., 2007; Murray et al., 2007; Squire et al., 2007). In humans, PrC has been implicated in recognition memory for many different stimulus classes, including objects, faces, and words (see Diana et al., 2007; Kim 2013). A noticeable but less consistent body of research also suggests that recognition of prior occurrence does not require the integrity of the hippocampus, and can proceed normally in the absence of recollection of contextual detail about specific past stimulus encounters (see Montaldi and Mayes, 2010; Ranganath, 2010; Wixted et al., 2010; Yonelinas et al., 2010, for review and discussion). Such acontextual item-based recognition has been linked to phenomenological feelings of familiarity. While numerous neuroimaging studies have focused on dissociations between perirhinal and hippocampal contributions to recognition memory, the precise nature of PrC computations and representations that support item-based recognition still remains poorly understood. Important outstanding questions concern how item-based familiarity is reflected in the functional magnetic resonance imaging (fMRI) blood-oxygen-level dependent (BOLD) response and whether pertinent signals in PrC show content specificity.

Electrophysiological evidence from a number of studies in rodents and in nonhuman primates suggests that the mechanism by which PrC could code for recognition of prior occurrence is a decrease in neuronal firing rate (i.e., repetition suppression; Zhu et al., 1995; Desimone, 1996; Ringo, 1996; Xiang and Brown, 1998; Aggleton et al., 2012). It has been reported that as many as 25% of PrC neurons in macaque monkeys show response decrements for familiar as compared to novel objects in the context of delayed-matching to sample, delayed nonmatching to sample, or continuous recognition-memory tasks (Brown et al., 1987; Sobotka

¹ Department of Psychology, The Brain and Mind Institute, Western University, London, Ontario, Canada; ² Department of Psychology, University of Toronto, Toronto, Ontario, Canada; ³ Department of Psychological and Brain Sciences, University of Massachusetts, Amherst, Massachusetts; ⁴ Department of Physiology and Pharmacology, Western University, London, Ontario, Canada; ⁵ Rotman Institute of Philosophy, Western University, London, Ontario, Canada; ⁶ Department of Philosophy, Western University, London, Ontario, Canada; ⁷ Baycrest Centre, Rotman Research Institute, Toronto, Ontario, Canada

Grant sponsor: Canadian Institutes of Health Research; Grant number: MOP93644.

*Correspondence to: Stefan Köhler, The Brain and Mind Institute, Natural Sciences Building, University of Western Ontario, London, ON, N6A 5C2 Canada. E-mail: stefank@uwo.ca

Accepted for publication 8 September 2015.

DOI 10.1002/hipo.22531

Published online 19 September 2015 in Wiley Online Library (wileyonlinelibrary.com).

and Ringo, 1996; Xiang and Brown, 1998). It should be noted, however, that not all studies with electrophysiological recordings have revealed response reductions in PrC in association with the repeated occurrence of objects. For example, Hölscher et al. (2003) reported that the response of some PrC neurons in rhesus monkeys showed a gradual increase in activity over hundreds of presentations of the same objects in the context of a delayed matching to sample task. Moreover, Thome et al. (2012) failed to observe any reduction in PrC firing rates for stimulus repetitions in a passive viewing experiment in rodents; the latter finding has been interpreted to suggest that such signals could be task dependent (Brown et al., 2012). Alternatively, this finding is open to the interpretation that item-based recognition may not be coded exclusively by a reduction in local firing rates but may involve distributed coding over populations of neurons in PrC (Thome et al., 2012; Burke et al., 2014). Evidence for such a coding schema has been reported in physiological recordings of activity in neighboring inferotemporal cortex, where it has been found that object identity can be decoded from population activity (Hung et al., 2005).

Findings obtained with fMRI in humans have also linked item-based recognition to differential PrC responses for previously studied as compared to novel items (see Diana, et al., 2007; Kim, 2013, for review). Notably, several studies have demonstrated a negative relationship between confidence in the perceived “oldness” of test items, which is often assumed to track item-memory strength, and the BOLD response in PrC (e.g., Gonsalves et al., 2005; Daselaar et al., 2006a,b; Montaldi et al., 2006; Wang et al., 2014). Given that this decrease in response for previously encountered items parallels the pattern of repetition suppression in electrophysiological recordings, it has attracted considerable attention in the literature. However, there are numerous challenges associated with mapping repetition effects in single-cell recordings onto fMRI BOLD signals (see Henson and Rugg, 2003; Grill-Spector et al., 2006; Gotts et al., 2012, for discussion). In fact, there is evidence to suggest that BOLD activity is more closely related to local-field potentials detected with multiunit recordings than to neuronal spiking (Logothetis et al., 2001; Logothetis, 2008; but, see Lee et al., 2010). Against this background, it is interesting that some fMRI studies have also revealed a relative increase in BOLD signal for familiar as compared to novel items in PrC (e.g., Kafkas and Montaldi, 2012), with some investigations even reporting increases and decreases in distinct PrC clusters in the same study (Yassa and Stark, 2008; Heusser et al., 2013). Such findings point to the possibility that item-based recognition-memory signals are reflected in spatially distributed patterns of decreases and increases in the fMRI BOLD responses across PrC. We note that stimulus repetition has also been linked to heterogeneous BOLD responses (i.e., decreases and increases) in other brain regions, including more posterior aspects of the ventral visual pathway (e.g., de Gardelle et al., 2013; see Segaert et al., 2013 for review).

Most studies on the role of PrC in recognition memory have examined the neural correlates of item recognition with univariate statistical analyses that probe for clusters of contiguous voxels with homogeneous response profiles (e.g., Daselaar et al., 2006a)

or by averaging activity across all voxels in anatomically defined regions of interest (e.g., Wang et al., 2014). Multivoxel pattern analyses (MVPA) of fMRI data, by contrast, can detect information carried in activity patterns distributed over multiple voxels even when these voxels are not part of a contiguous cluster, and, critically, even when they show a heterogeneous directional response to an experimental manipulation such as stimulus repetition (e.g., de Gardelle et al., 2013; for review, see Norman et al., 2006; Rissman and Wagner, 2012; Tong and Pratte, 2012). In an MVPA-based fMRI study, we recently revealed patterns of BOLD activity in MTL structures that can distinguish between subjectively familiar and novel items (Martin et al., 2013). Participants were scanned while they made recognition-memory judgments for visual stimuli from several different object categories. Using a linear support vector machine (SVM) with separate training for items from each stimulus category, we were able to successfully classify the perceived familiarity of individual faces and chairs from activity patterns in right PrC.

For the present report, we took advantage of the dataset collected by Martin et al. (2013) and conducted new analyses that aimed to address precisely how PrC may code for the prior occurrence of visual stimuli. Our primary goal was to determine whether patterns of BOLD responses in PrC that allow for the classification of the perceived familiarity of visual stimuli indeed comprised voxels with heterogeneous response profiles in terms of direction, and, if so, whether classification could still be successful if patterns were restricted to include only voxels with changes in one direction. To obtain further leverage on the issue, we also investigated whether differences in classification accuracy for any such voxel patterns are correlated with differences in behavioral performance across participants.

Insight as to how PrC codes for the prior occurrence of visual stimuli can also be gained from considering whether patterns of PrC activity that carry information about familiarity show specificity in their response for particular stimulus categories (Cowell et al., 2010; Graham et al., 2010). Inasmuch as item-based recognition-memory signals are, by their definition, based on the object itself rather than any contextual information about a pertinent prior encounter, the category to which an object belongs may play an important role in their neural organization (Martin et al., 2012, for further discussion). Prior fMRI research that has examined category specific responses for visually presented stimuli in nonmnemonic tasks has revealed two types of effects in more posterior occipitotemporal regions. First, studies based on univariate analyses have revealed clusters of contiguous voxels in the ventral visual pathway that show maximal responses for exemplars from a specific visual category. Such clusters are often referred to as category-selective regions and have been reported for several ecologically relevant categories, including faces, body parts, words, and places (see Op de Beeck et al., 2008, for review). Second, MVPA-based studies have revealed patterns of activity distributed across voxels in wider swaths of posterior ventral temporal cortex that show category-specific responses even when clusters with preferential responses are excluded, and even for categories that are not associated with any contiguous clusters in univariate analyses (see Grill-Spector and Weiner, 2014;

Haxby et al., 2014, for review). For example, the latter types of studies have revealed the presence of informational content relevant for face identity in distributed patterns of activity in ventral visual pathway regions that go beyond the classic lateral occipital and posterior fusiform “face areas” previously identified in univariate analyses (Haxby et al., 2001; Kriegeskorte et al., 2008; Nestor et al., 2011). At present, it remains unknown whether patterns of recognition-memory signals for faces in PrC show a similarly distributed functional organization or are found, instead, in a more circumscribed area. This issue is of particular interest given that a category-specific region with a preferential response for faces has recently been identified in the anterior collateral sulcus of PrC, a region sometimes referred to as the human homologue of the anterior temporal face patch (ATFP), which is situated more laterally in nonhuman primates (Rajimehr et al., 2009; Nestor et al., 2011; Nasr and Tootell, 2012; Rossion et al., 2012; O’Neil et al., 2013, 2014; Collins and Olson, 2014). As is the case for other nodes that form the face processing network in humans (Behrmann and Plaut, 2013), this region has most consistently been identified in the right hemisphere.

In this study, we examined whether voxels in right PrC with diagnostic relevance for classification of recognition-memory judgments for faces show category specificity when probed with the type of functional localizer that has been used to identify face-specific responses in the ventral visual pathway under passive-viewing conditions in many prior studies. We specifically focused on faces because they are the category most widely studied in the broader literature on category specificity in human brain responses, and because an emerging body of research implicates aspects of PrC (i.e., the ATFP) as part of the face processing network. We pursued these issues within the context of our broader goal of characterizing the nature of item-based recognition-memory signals in human PrC, which, as discussed, also addressed whether these signals are reflected in distributed patterns of voxels that show decreases and increases in response to prior exposure.

MATERIALS AND METHODS

Participants

Nineteen right-handed individuals participated in the study (21–30 years of age, mean age = 25.2 years; 12 females). All participants were screened for the absence of a history of neurological disorders. Data from one participant were excluded from all analyses due to excessive head movement during scanning. Participants received financial compensation for their participation and provided informed consent according to procedures approved by the University of Western Ontario Health Sciences Research Ethics Board.

Stimuli and Behavioral Procedure

Stimuli were grayscale images depicting exemplars from three different object categories (i.e., faces, chairs, and buildings), although this study focused only on patterns of BOLD activity

related to the familiarity of face trials. Each target object was presented in isolation on a homogeneous, white background. The size of each image was limited to be 375×250 pixels, with at least one dimension corresponding to these limits. For the purpose of counterbalancing, images from each stimulus category were divided into three sets of 40 objects, two of which (i.e., 80 objects) served as items presented prior to scanning during a study session, and as corresponding targets in the scanned recognition-memory test stage. The remaining 40 items served as novel lures in the recognition task. Assignment of item sets to either target or lure lists was counterbalanced across participants.

The experimental task consisted of two discrete stages: an encoding session and a subsequently scanned recognition-memory test. The initial encoding session was separated into six blocked sequences that were counterbalanced across participants. Each block consisted of 40 trials corresponding to one target list. Stimuli were presented for 3,000 ms with a 2,000 ms fixation inter-stimulus interval, and participants were asked to rate the relative attractiveness of each face, comfort of each chair, or value of each building using a five-point scale.

Participants subsequently completed a scanned recognition-memory test consisting of 80 previously studied targets and 40 lures from each category, for a total of 360 trials distributed over eight functional runs of equal length and composition. Of these trials, 120 corresponded to presentation of face stimuli (i.e., 80 studied and 40 novel lures). Stimuli were presented for 2,500 ms, with a jittered fixation-baseline separating trials (jitter sequence was optimized using the OptSeq2 algorithm; <http://surfer.nmr.mgh.harvard.edu/optseq/>). While in the scanner, participants viewed the stimulus display through a mirror at a distance that yielded an approximate object size of $18 \times 13^\circ$ visual angle. For their recognition judgments, participants were instructed to provide a rating of perceived familiarity on a scale between one (least familiar) and four (most familiar), with a fifth response option corresponding to recollection. Critically, they were asked to respond with a fast and intuitive assessment of their perceived item familiarity and to avoid attempting to recollect contextual details from the encoding stage of the experiment (see Dobbins and Han, 2006; Montaldi et al., 2006; Quamme et al., 2010; Martin et al., 2013, for further discussion). However, they could indicate recollection if it occurred unintentionally. Recollection of contextual details was defined as any situation that involved conscious awareness of information about the past item encounter that was independent of perceptual details of the stimulus itself, such as internal thoughts and associations that were formed during the initial item encounter.

Functional Localizer Tasks

Subsequent to the experimental task, each participant completed two independent functional localizer scans (which were not considered in our initial report, Martin et al., 2013). The localizer task followed a protocol that has previously been used in several other studies from our lab (e.g., Ganel et al., 2006; O’Neil et al., 2009; Cate et al., 2011; O’Neil et al., 2013) and is similar to that used in many other studies in the visual neuroscience literature more broadly. It involved presentation of

grayscale faces, common objects, and places (scene landscapes and buildings with naturally occurring background) under passive viewing instructions. Importantly, stimuli employed in the localizer task were different from those comprising the experimental recognition-memory task. Stimuli from each category were presented in a blocked manner with alternating blocks of scrambled images corresponding to each stimulus category.

fMRI Acquisition Protocol

All magnetic resonance imaging (MRI) data were acquired on a Siemens TIM Trio 3-Tesla scanner with a high-resolution fMRI protocol optimized for MTL examination. Functional MRI volumes were collected using a T2*-weighted single-shot gradient-echo-planar acquisition sequence (repetition time (TR) = 2,500 ms, echo time (TE) = 26 ms, slice thickness = 2 mm, in-plane resolution = 2 × 2 mm, field of view (FOV) = 220 mm × 220 mm, flip angle = 90°). Each functional volume included 37 contiguous slices collected in an interleaved manner. For each experimental run ($n = 8$ per participant), 176 volumes were collected. Each localizer scan ($n = 2$ per participant) consisted of 150 functional volumes. To optimize MR signal in the anterior temporal lobes, a transverse orientation was chosen with the effort to include the entire temporal lobes and as much visual cortex as possible. This slice selection resulted in full coverage of the ventral aspects of occipital and full coverage of the entire temporal lobes in all participants, with exclusion of the most dorsal aspects of frontal and parietal cortices, as well as occipital cortex in some participants. A saturation band was applied during functional runs to minimize artifacts related to eye-movements and the sinus cavity. T1-weighted anatomical images were obtained using an ADNI MPRAGE sequence (192 slices, TR = 2,300 ms, TE = 4.25 ms, 1 mm isotropic voxels, FOV = 240 × 256 mm, flip angle = 9°).

fMRI Data Preprocessing

fMRI data were preprocessed in native space using BrainVoyager QX version 2.3 (Brain Innovation). Functional images were slice-scan time corrected, 3-D motion corrected with reference to the functional volume taken just prior to the anatomical scan, and high-pass filtered using a Fourier basis set of 2 cycles/run (including linear trend). Images were then coregistered with the anatomical image and aligned with the anterior commissure–posterior commissure (AC–PC) plane. For the purpose of MVPA, experimental data were minimally smoothed using a 3-D Gaussian kernel with a full-width at half maximum of 3 mm. Functional data from the localizer experiment were smoothed using a 3-D Gaussian kernel with a full-width at half maximum of 8 mm. Functional data were convolved using a standard double gamma hemodynamic response function (Friston, 1998). Participant-specific general linear models (GLM) of experimental data allowed for extraction of z -scored trial-specific beta estimates in all voxels of interest. Beta estimates derived from a modeled HRF were chosen as target measure for the MVPA (i.e., as classifier input) because they are particularly well suited to account for overlap in the hemodynamic response in fast-event-related designs (Misaki et al., 2010). Changes in mean

intensity across runs were modeled by including them as predictor of no interest in the participant-specific GLMs.

Identification of the AFTP in Functional Localizer Scans

This region of interest was defined functionally, for individual participants, in the right hemisphere, using data from the localizer scans. A general linear model was specified for each localizer run with faces, scenes, and objects as predictors. Scrambled images served as the baseline condition. Participants' AFTPs were derived from the contrast (faces > scenes) in a subject-specific fixed-effects model. ROIs were defined based on activation maps that were thresholded using a liberal statistical criterion (uncorrected P -value < 0.05), respecting anatomical boundaries for PrC as described by Pruessner et al. (2000, 2002).

Anatomical Definition of PrC for MVPA of Recognition Signals in Experimental Task

To conduct MVPA, an anatomically defined ROI for PrC in the right hemisphere was created in native MRI space with manual tracing separately for each participant, again using anatomical landmarks for demarcation of PrC from surrounding cortical structures in the MTL as described by Pruessner et al. (2000, 2002).

MVPA of fMRI Data

For the purpose of classification, data were collapsed across response options such that the familiar class of face trials corresponded to hits and false alarms at the two highest levels of familiarity (i.e., response options 3 and 4) and the novel class of face trials comprised hits and false alarms at the two lowest levels of familiarity (i.e., response options 1 and 2). To eliminate potential classifier bias related to unequal trial numbers we used a pseudorandom sampling procedure that equated the number of trials between the familiar and novel classes. This procedure was repeated over 10 iterations to ensure that all trials for a given participant were included in the classification analysis at least one time. Accordingly, 10 separate instances of the classification analysis (i.e., cross-validated classifier training and testing) were completed and inferential statistical analyses were performed on classifier accuracy averaged over these 10 iterations. Across participants, the average number of trials included for the classification of faces at each familiarity level (i.e., familiar vs. novel) was 39.8 (range = 24–56).

Pattern classification analyses were conducted using the Princeton MVPA toolbox (<http://www.pni.princeton.edu/mvpa>) and custom MATLAB code (The MathWorks, Natick, MA). As a first step, we performed feature selection in order to minimize the influence of noise in the functional data. The feature selection procedure used here allowed for multivariate classification of perceived familiarity of faces based on activity within a subset of PrC voxels that were not necessarily clustered in any systematic manner and showed either homogeneous or heterogeneous response profiles. Specifically, feature selection was based on voxel-wise measures of discriminability (i.e., t -tests between familiar and novel). When contrasted

TABLE 1.

Recognition-Response Distribution and Discrimination Estimate for Faces

	Percentage responses to novel items					Percentage responses to studied items					Discrimination Familiarity d'
	1	2	3	4	R	1	2	3	4	R	
Mean	42.9	30.9	16.3	8.0	2.0	21.9	26.7	20.8	19.4	11.2	0.64
SEM	5.1	3.1	1.8	1.9	0.9	3.1	2.1	1.4	2.1	2.0	0.08

For behavioral signal-detection and fMRI analyses, novel responses correspond to 1 and 2 collapsed, with familiar responses corresponding to 3 and 4 collapsed. R corresponds to recollection trials not considered in the current analyses.

with multivariate feature selection procedures, such as a multivariate searchlight which considers weighted combinations of voxel responses for class separation, the primary advantage of the current approach pertains to increased sensitivity for detection of cognitive states coded in activity patterns comprised of spatially distributed voxels. Accordingly, this approach is sensitive to meaningful patterns that are distributed beyond the spatial scale of a searchlight.

Feature selection was performed in each participant separately by choosing the subset of voxels in right PrC that appeared most informative for classification based on an initial univariate statistical analysis (see Norman et al., 2006, for discussion). Specifically, a t -test was conducted between beta values for familiar and novel trials in all voxels in right PrC for SVM training data for each cross-validation separately. All voxels were subsequently rank ordered according to their obtained t -statistic and those corresponding to the top 10% of that ranking were selected as features included for SVM classification (see below for additional detail regarding directional and nondirectional feature selection procedures). For all familiar-novel classifications, this analysis was also performed separately for each of the 10 iterations of item sampling. This feature selection procedure yielded an average of 55.1 functional voxels ($2 \times 2 \times 2$ mm) in right PrC across participants.

A linear SVM (libSVM, <http://www.csie.ntu.edu.tw/~cjlin/libsvm>) was used for classification of beta values with a linear kernel function and a constant cost parameter of $C = 1$. For each cross-validated classification analysis, the SVM was trained on all but two face trials; those trials not included in the training dataset (i.e., one familiar and one novel trial) subsequently served as test trials for assessment of classifier performance. This train and test procedure was completed in a fully cross-validated manner such that every trial served as the test stimulus for classification. For each trial in the test set, the classifier returned a probability estimate that indicated the likelihood that the activity pattern corresponded to either the familiar or novel class that was used for SVM training purposes. Probability estimates were then binarized in a winner takes all manner; classification was either correct (i.e., when the “true” experimental condition was assigned the highest probability) or incorrect. Averaged across all 10 iterations, classifier accuracy for the perceived familiarity of faces reflects the percentage of test trials that were classified correctly in this binary manner. To obtain inferential statistics, we examined whether average classification performance was above chance

(i.e., >0.5). For this purpose, we used a single sample t -test to test against a population mean of chance level.

RESULTS

As summarized in our initial report (Martin et al., 2013), the behavioral accuracy of item-based recognition decisions was quantified using d' based on response options 1–4 (i.e., least to most familiar), and corrected assuming independence between familiarity and recollection (Yonelinas, 1999). Although performance levels were low, as expected due to the high similarity between all faces, item-based discrimination was above chance ($d' = 0.64$, $t_{17} > 4.81$, $P < 0.001$). We note that behavioral data did not allow us to constrain fMRI analyses to more specific response options (e.g., level 1 vs. level 4, rather than levels 1 and 2 vs. 3 and 4) due to restrictions in item numbers in participants' response distributions. Similarly, the limited number of R responses (collapsed across hits and false alarms $M = 9.56$) did not allow for any investigation of associated fMRI responses. The observed distribution of recognition responses is presented in Table 1.

As previously reported, MVPA-based analyses of right PrC activity allowed us to successfully decode the perceived familiarity or novelty of individual faces with a mean classifier accuracy of 57% (Bonferroni corrected $P < 0.001$), and 14/18 participants showed activity patterns that could be classified with numerical above chance performance (i.e., $>50\%$). By contrast, classifier accuracy in left PrC was not significantly different from chance performance level (51%). Thus, we specifically focused only on right PrC in all further analyses.

Direction of Signal Change in PrC Activity Patterns That Allow for Decoding of Recognition-Memory Decisions for Faces

Successful decoding of recognition-memory decisions, as summarized above, indicates greater within- than between-class similarity in PrC activity patterns for subjectively familiar as compared to novel trials. To characterize precisely how these class differences are reflected in BOLD activity we first examined the extent to which voxels with diagnostic relevance for classification showed the same or a varied response with respect

Voxel-Wise Directional Effects and Overlap with Anterior Temporal Face Patch

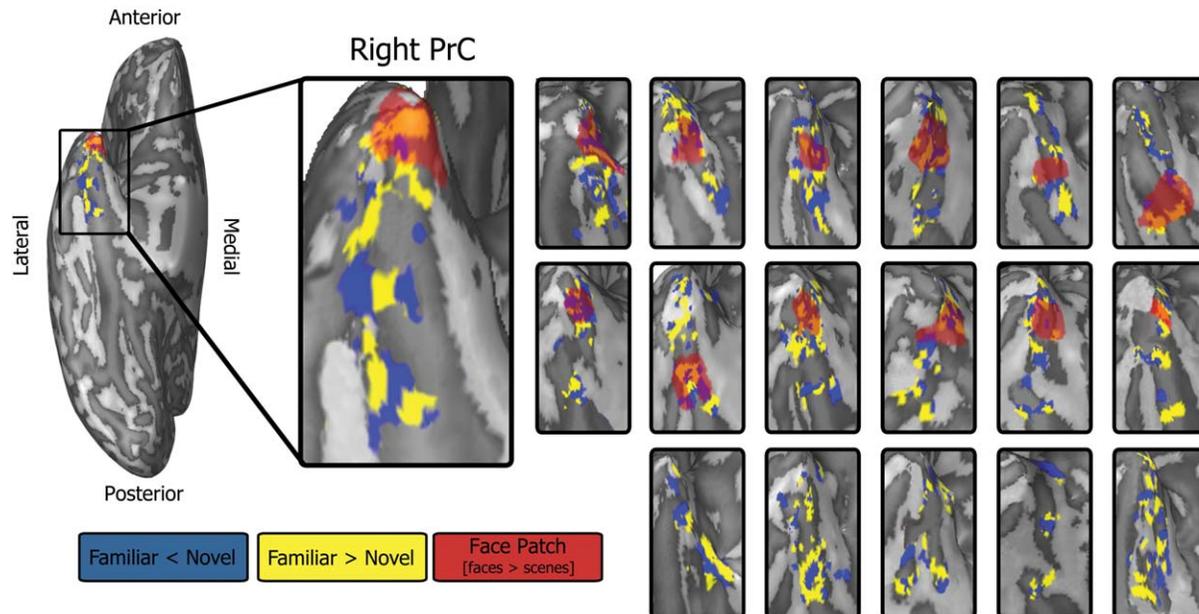


FIGURE 1. Spatial distribution of voxels with diagnostic relevance for decoding of item-recognition decisions in each participant. For illustrative purposes, the data presented for each participant refer to one representative, fully cross-validated iteration of the classification analysis. Blue coloring denotes voxels that showed response decrements to familiar as compared to novel

faces. Yellow coloring denotes voxels that showed response increments to familiar as compared to novel faces. Red color patches correspond to the ATRP in those participants for whom one could be identified using independent functional localizer data (faces > scenes). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

to familiarity in terms of direction. Decoding results from our previous investigation (Martin et al., 2013) were obtained following a commonly used nondirectional feature selection procedure that was based on initial voxelwise measures of experimental effects in a GLM-derived test statistic (i.e., t -values for contrast between familiar and novel). Specifically, voxels were rank ordered according to the absolute value of their obtained t -statistic (without consideration of any statistical threshold), and the top 10% of voxels were selected for the purpose of SVM training and classification. With the feature selection described, voxels in which activity decreased with familiarity, as well as others in which activity increased with familiarity could be included in the resulting patterns. The spatial distribution of voxels with diagnostic relevance for the classification of face recognition decisions is displayed in Figure 1, color-coded for direction of change.

It is possible that the diagnostic voxels selected and depicted in Figure 1 showed a difference in mean (when averaged across voxels), and that successful decoding reflected SVM leveraging of this spatially distributed mean difference (see Coutanche, 2013, for further discussion). To evaluate this possibility, we calculated mean beta values for subjectively familiar and novel trials based on activity in voxels that survived feature selection in the majority of classification analysis iterations (i.e., included in a minimum of 6 out of 10 trial sampling iterations). Mean

beta values, collapsed across participants, are presented in the left-most bars of Figure 2. At the group level, no significant difference was observed between beta values that corresponded to familiar and novel trials ($t_{17} = 0.21$, $P = 0.83$, $d = 0.10$). Of the voxels included, only 46% showed a numerical decrease in activation for familiar trials, indicating that both response directions were strongly represented in the selected voxel populations. At the single subject level, only 5 of 18 participants had a significant mean difference ($P < 0.05$) between both types of trials; moreover, only two of these five participants showed a decrease (familiar < novel) in beta values when averaged across the selected PrC voxels. Although these mean differences between classes were clearly limited and not consistent across participants, we also sought to determine whether classification would still be successful after demeaning familiar and novel beta values. Specifically, in this analysis, beta values across all voxels that survived feature selection were z -scored for each trial and participant separately; this ensures that mean differences are exactly zero. Critically, we found that classifier performance remained above chance in this scenario ($M = 56%$, $t_{17} = 5.04$, $P < 0.001$, $d = 2.44$). Decoding results obtained with and without z -scored beta values are presented in the left-most bars of Figure 3 for comparison. These results suggest that successful decoding of recognition-memory decisions does not rely on the presence of a mean difference. In turn, they

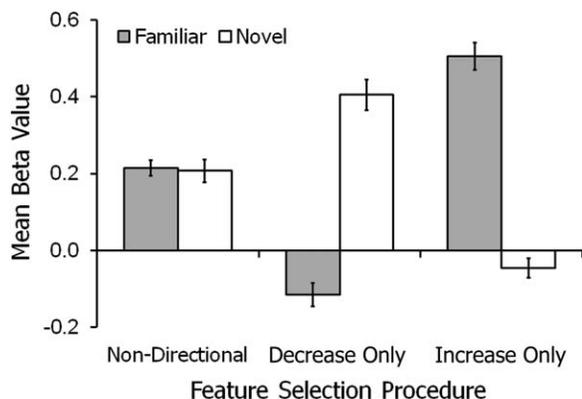


FIGURE 2. Mean beta values in right PrC for familiar and novel trials. Mean beta values were calculated across participants based on voxels that were reliably selected for classification following nondirectional and directional feature selection procedures. Reliable voxels were those that survived feature selection in at least 6 out of 10 analysis iterations. All error bars indicate the SEM calculated across participants.

suggest that feature selection that is blind to direction yields successful classification of recognition decisions based on patterns of voxels that have heterogeneous response profiles.

Decoding of Recognition-Memory Decisions from PrC Activity Patterns When Direction of Signal Change Is Constrained

We next sought to determine whether successful classification of recognition decisions necessitates consideration of voxels with heterogeneous response profiles. We addressed this question using MVPA based on a feature selection approach that allowed for inclusion of voxels with a change in signal in only one direction. Toward this end, we ran two separate analyses with feature selection constrained to be based on voxels with decreases or increases in signal, respectively. Voxels were ranked according to raw, rather than absolute, t -values and those corresponding to the top or bottom 10% of these rankings were selected for the two separate MVPAs. Thus, in the first set, all voxels showed a decrease in response for familiar as compared to novel trials, while voxels in the second set showed the opposite response profile.

Mean beta values for voxels that survived this directionally constrained feature selection are presented in the center (familiar < novel; i.e., decrease only) and right-most (familiar > novel; i.e., increase only) bars in Figure 2 collapsed across participants. Notably, directionally constrained feature selection still resulted in classifier accuracy that was significantly greater than chance for both types of analyses (see Fig. 3; decrease only $M = 63\%$, $t_{17} = 11.67$, $P < 0.001$, $d = 5.66$; increase only $M = 63\%$, $t_{17} = 10.32$, $P < 0.001$, $d = 5.01$). These results suggest that information pertaining to item recognition decisions can also be successfully decoded from patterns of voxels with a homogeneous response profile between familiar and novel trials, regardless of whether this difference reflects a decrease or an increase.

Relationship Between Classifier Accuracy for PrC Activity Patterns and Behavioral Recognition-Memory Performance

Taken together, the results of the analyses presented thus far suggest that recognition decisions can be successfully decoded from distributed activity patterns in PrC regardless of whether signal changes are constrained in terms of direction or not. Is it possible to determine which of these different patterns in PrC is most relevant for successful behavior, that is, memory accuracy? Note that all analyses presented involve classification of recognition decisions without taking their accuracy, on a trial-by-trial basis, into account. This approach was chosen so as to maximize the number of trials available for training of the classifier. To get leverage in answering questions about behavioral performance, however, one can also examine the relationship between classifier performance and behavioral accuracy on a subject-by-subject basis (i.e., by focusing on interindividual differences). Accordingly, we assessed this relationship for each of the three feature selection procedures previously described. To obtain an estimate of behavioral performance, we calculated familiarity-based discrimination between targets and lures using a measure derived from signal-detection theory (d'). The results of these correlation analyses are plotted in Figure 4. Critically, we found a significant positive correlation ($r = 0.47$, $P < 0.05$) between familiarity-based discrimination and classifier performance for voxel patterns in PrC using the unconstrained feature-selection procedure that allowed for the inclusion of voxels with decreases or increases in their response. In other words, those participants in whom classification of recognition decisions from patterns of activity in PrC was more successful tended to perform better in familiarity-based discrimination of faces. By contrast, classifier accuracy and behavioral performance were not significantly correlated when feature selection was constrained to include only voxels with changes in signal in one direction (for voxels showing familiar < novel: $r = 0.16$, $P = 0.27$; for voxels showing

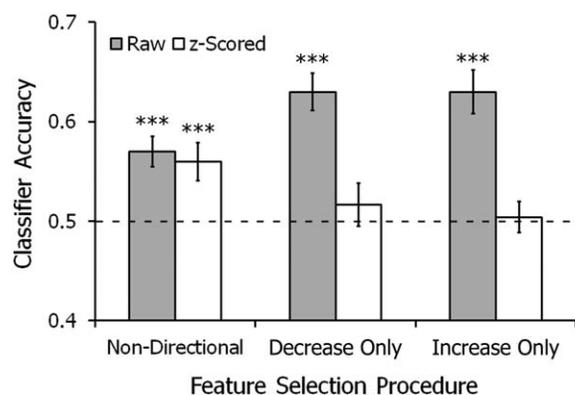


FIGURE 3. Decoding accuracy for perceived familiar versus novel trials from raw and z-scored patterns of activation across feature selection approaches. z-scoring was performed on beta values across all voxels for each trial and participant separately to ensure that familiar and novel trials were equated at the level of mean activation. Dashed line indicates chance level for classification. All error bars indicate the SEM calculated across participants. *** $P < 0.001$.

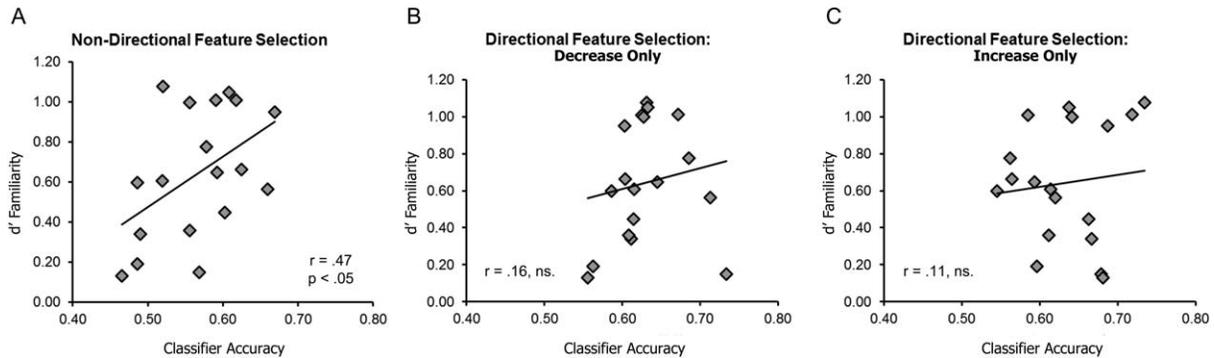


FIGURE 4. Pearson correlations between classifier accuracy and familiarity-based behavioral discrimination (d') across participants. (A) Correlation obtained following nondirectional feature selection. (B) Correlation obtained following directional feature selection of voxels that showed activity reductions for familiar relative to novel trials. (C) Correlation obtained following directional feature selection of voxels that showed increased activity for familiar relative to novel trials.

familiar > novel: $r = 0.11$, $P = 0.33$). These data suggest that, although successful decoding of recognition-memory decisions from activity patterns in PrC can be obtained in multiple ways, only decoding with patterns that consist of voxels with increases and decreases in signal shows a relationship to behavioral memory performance.

The ATFP and Item-Based Recognition-Memory Signals for Faces

To characterize the relationship between recognition-memory signals for faces in PrC and the ATFP, we first assessed the extent to which voxels with diagnostic relevance for the classification of recognition-memory decision for faces in PrC overlapped with the ATFP. For this purpose, we concentrated on the MVPA approach with a feature-selection procedure that allowed for inclusion of voxels with either direction of signal change. Of the right PrC voxels that were selected in the majority of iterations for successful classification of recognition decisions, 17.7% (averaged across participants; range = 10.4–25.3%) overlapped with the ATFP in those participants for whom it could be defined (13 of 18 participants in the sample). This overlap is illustrated in Figure 1. Thus, the majority of voxels that were part of the patterns that allowed for successful classification of recognition-memory decisions for faces were located outside of the ATFP, even when the latter was defined at the individual subject level. Next, we performed MVPA to determine whether recognition decisions for faces could be decoded from activity restricted to voxels comprising the ATFP in each participant, without any additional feature selection. Notably, even though the average number of voxels that entered these analyses was comparable ($M = 49.7$) to that which entered our MVPA that were based on GLM-derived feature selection without any constraints in spatial distribution ($M = 55.1$), classifier performance did not exceed chance level with this approach ($M = 51\%$, $t_{12} = 0.87$, $P < 0.20$, $d = 0.42$). In a third step, we also conducted a complementary MVPA excluding PrC voxels that were part of the contiguous clusters that defined the ATFP in these 13 participants. Critically, this analysis revealed

above chance classifier performance ($M = 57\%$, $t_{12} = 2.29$, $P < 0.05$, $d = 1.11$). Taken together, the results from these analyses converge to suggest that voxels carrying information pertinent to recognition-memory decisions for faces are spatially distributed in PrC, and not confined to the ATFP.

Category Specificity of Responses in PrC Voxel Patterns That Allow for Classification of Recognition-Memory Decisions for Faces

Although the results of the previous analyses suggest that face familiarity signals are spatially distributed and extend beyond the ATFP, this finding does not refute the notion that these patterns of distributed activity may code for familiarity in a category-specific manner. To address the issue of category specificity, we first examined whether PrC voxels that form the patterns allowing for classification of recognition decisions for faces show a preferential response to face stimuli when probed with an independent functional localizer. Specifically, we examined localizer activity in voxels that were consistently included in feature selection in at least 6 of our 10 classification iterations for decoding of recognition-memory decisions when no constraints in signal direction were imposed. A histogram of mean difference scores (averaged across blocks) for faces as compared to common objects, and faces as compared to scenes, are shown in Figure 5. Notably, the distribution is visually skewed toward positive values, hinting at predominant preferential tuning for faces in these voxel populations. Statistically, the mean difference score was indeed different from zero in both comparisons with other stimulus categories (faces > objects, $M = 0.19$, $t_{17} = 3.35$, $P < 0.01$, $d = 1.62$; faces > scenes, $M = 0.28$, $t_{17} = 4.95$, $P < 0.001$, $d = 2.40$). A comparable pattern of results was obtained even after exclusion of voxels that overlapped with the ATFP (faces > objects, $M = 0.12$, $t_{17} = 2.31$, $P < 0.05$, $d = 1.12$; faces > places, $M = 0.16$, $t_{17} = 3.09$, $P < 0.01$, $d = 1.49$). These data suggest that the voxels in patterns of PrC activity that allow for decoding of recognition-memory decisions for faces do indeed show a modest tuning preference for faces in a classic functional-localizer paradigm.

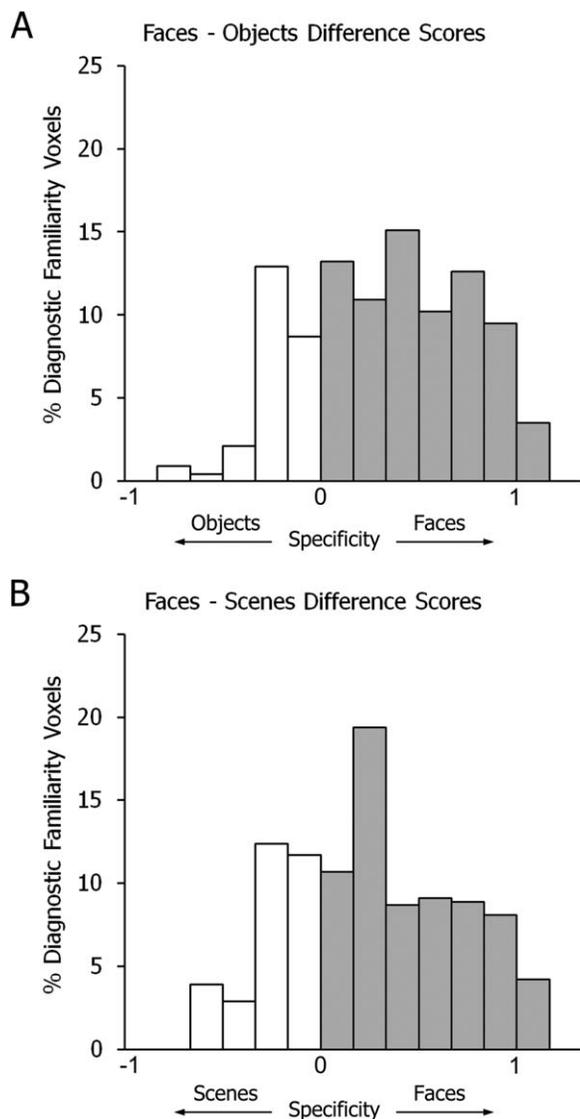


FIGURE 5. Category preference revealed with functional localizer data in right PrC voxels with diagnostic relevance for decoding of face recognition-memory decisions. Histograms depict the proportion of voxels that show a preference for either (A) faces (shaded bars) or objects (open bars), and (B) faces (shaded bars) or scenes (open bars). Difference scores were calculated based on activity from the functional localizer scans in voxels with diagnostic relevance for decoding of recognition decisions for each participant separately (faces—objects; faces—scenes). These values were then collapsed across participants and plotted as a percentage of the total number of voxels.

Distributed Category-Specific Localizer Responses and Item-Based Recognition-Memory Signals for Faces

In a third set of analyses that addressed category-specificity, we examined whether voxels that showed preferential responses to faces in the passive viewing localizer task could be used to decode face recognition decisions from our experimental data. Thus, we first used MVPA to identify voxels in PrC that show

a preference for face stimuli, without any requirement to be part of a contiguous voxel clusters, and after exclusion of the ATFP (in those participants for whom one could be identified). For feature selection, we selected the top 10% of voxels that were preferentially tuned to faces in the two critical contrasts of our localizer data, that is, (faces > objects) and (faces > scenes), respectively. Using this approach, we were able to classify stimulus category robustly, with above chance accuracy, in both comparisons (faces vs. objects: $M = 0.74$, $t_{17} = 7.51$, $P < 0.001$, $d = 3.64$; faces vs. scenes: $M = 0.79$, $t_{17} = 8.22$, $P < 0.001$, $d = 3.99$).

Importantly, the distributed voxel patterns that yielded above chance classification of stimulus category in our localizer data overlapped considerably with those that allowed for decoding of face recognition decisions in our experimental task (when examined with nondirectional feature selection). The overlap was 68% and 70%, respectively. Given this substantial overlap, we also evaluated whether we could directly decode face recognition decisions in the experimental task employing the MVPA-based localizer data for initial feature selection. Critically, we were successful with this approach; voxels that allowed for the decoding of faces versus objects in our localizer data could also be used to classify face familiarity ($M = 0.58$, $t_{17} = 3.17$, $P < 0.001$, $d = 1.54$). We obtained a similar result when we focused on voxels that allowed for classification of faces versus scenes ($M = 0.59$, $t_{17} = 3.09$, $P < 0.001$, $d = 1.50$). Overall, the results of the different sets of analyses we used to compare activity in the functional localizer task and the experimental task converge to suggest category specificity in item-based recognition-memory signals in PrC that is spatially distributed beyond the ATFP.

DISCUSSION

We used fMRI-based MVPA to characterize the patterns of activity in right PrC that allow for successful classification of item-based recognition-memory decisions for faces (i.e., perceived face familiarity). When no constraints for the direction of signal change were imposed in feature selection, these patterns included voxels with decreases as well as voxels with increases in signal. Moreover, successful classification was obtained in the absence of a mean difference in activity across the entire set of voxels in the patterns. While we also found above chance classification when analyses were constrained to include only voxels with signal changes in one direction, decoding accuracy was related to memory performance only when decreases and increases in signal were considered. At another level, comparison with data from an independent functional localizer revealed that the patterns in PrC that allowed for classification of the perceived familiarity of faces comprised voxels that show category specificity in their response. These voxels were found to be spatially distributed, and extended beyond the ATFP.

Extant fMRI evidence obtained with univariate statistical analyses has typically linked item-based familiarity to mean activity differences in clusters of contiguous PrC voxels with the same direction of signal change. Such effects have most often been associated with relative decreases in activity for old as compared to novel stimuli at the time of retrieval (Henson et al., 2003; Gonsalves et al., 2005; Daselaar et al., 2006a,b; Montaldi et al., 2006; Danckert et al., 2007; Wang et al., 2014), a finding that has attracted considerable attention due to its parallels in neurophysiological recordings in PrC. However, some fMRI studies have also reported relative increases in activity, or both types of effects for different clusters in the same study (e.g., Yassa and Stark, 2008; Kafkas and Montaldi, 2012; Heusser et al., 2013). Against this background, analytical approaches that focus on patterns of increases and decreases can offer a clear advantage for probing the role of PrC in memory processing. Indeed, while above chance classification of recognition decisions could still be observed in the current study when feature selection in MVPA was constrained to include only PrC voxels with decreases in signal, classification accuracy was related to the accuracy of corresponding behavior only when patterns included voxels with decreases as well as voxels with increases in their response.

That we revealed a relationship to memory performance only when increases and decreases in PrC signal were considered together, in the absence of significant mean-activation differences across all voxels in the patterns, suggests that these patterns could carry behaviorally relevant information over and above what individual voxels carry. Put another way, our data suggest that the pattern context in which voxels that show a response to prior exposure are embedded is of functional relevance for item-based recognition-memory decisions. Accordingly, these data can be seen as support for the notion that recognition-memory signals in PrC involve a distributed code that is carried over populations of elements, as has been suggested in the neurophysiological literature (Thome et al., 2012; Burke et al., 2014; see also Hung et al., 2005, for object coding in inferotemporal cortex). From a computational perspective, two models of recognition memory in PrC (or in MTL neocortex, including PrC) have also shown that it is possible to simulate stimulus familiarity via a mechanism in which stimulus representations become “sharpened” during encoding, manifested as an increase in the activation of some neural units and a relative decrease in the activation of others (Norman and O’Reilly, 2003; Cowell et al., 2006). Of course, even though the current fMRI data were obtained at very high resolution (voxel size of 2 mm in each dimension), every voxel would still comprise hundreds of thousands of neurons. As long as the mapping of signals from neurons to voxels remains incompletely understood, any fMRI evidence must ultimately be considered indirect as to the code used for neuronal signaling.

We also caution that the present fMRI results do not permit for inferences as to the dimensionality of the signal carried by voxel patterns of PrC activity that allow for decoding of item familiarity. A multidimensional code is typically characterized

with respect to multiple psychological or behavioral dimensions that are reflected in neural response patterns (see Diedrichsen et al., 2013; Davis et al., 2014, for discussion). Further research is required to specify these dimensions in relation to item-based recognition decisions, to probe how they map onto increases and decreases in signal, and to relate them to specific sources of stimulus information that could guide these memory decisions (Henson and Gagnepain, 2010).

A second question about item familiarity signals in PrC that was addressed in this study was whether diagnostic voxels in the activity patterns that allow for successful decoding are tuned to the specific stimulus category to which a judged object belongs. Results from several analyses we present, which compared data from the experimental task with data from an independent classic functional localizer of the kind typically used in research on category-specific responses in more posterior temporal and occipitotemporal region, suggest that this is indeed the case. Most notably, those PrC voxels that constituted the patterns that allowed for decoding of the perceived familiarity of faces responded preferentially to faces as compared to common objects, and faces as compared to scenes, when probed with the independent functional localizer (Fig. 5). Furthermore, we were also able to decode the perceived familiarity of faces from those voxels in PrC that constituted patterns with face-specific responses as identified with MVPA-based analyses of our functional localizer. Results from both types of analyses converge to suggest that recognition-memory signals in PrC adhere to a category-specific organization.

A more subtle aspect of our findings on category specificity is that we were able to decode the perceived familiarity of faces from widely distributed voxel patterns in PrC with faces-specific responses, but not when examination of voxel patterns was restricted to the AITFP. It is interesting to note the parallel between this finding and those in more posterior aspects of the ventral visual stream, where it has been found that spatially distributed voxel patterns in inferotemporal and occipitotemporal cortex carry information about faces even after exclusion of the fusiform face area (Haxby et al., 2001; Kriegeskorte et al., 2008). Our results suggest that even though this general representational structure in the ventral visual pathway (i.e., category specificity in contiguous local clusters and in spatially distributed patterns) is preserved in PrC, corresponding recognition-memory signals appear to be reflected primarily in the spatially distributed patterns.

Most prior fMRI literature on category-specific effects has focused on functional differences between distinct MTL structures, such as between PrC and the hippocampus (Lee et al., 2006; Barense et al., 2010; Graham et al., 2010), or between PrC and parahippocampal cortex (Staresina et al., 2011; Martin et al., 2013). Some research, however, has also revealed evidence for category specificity in distributed response patterns within PrC. For example, Liang et al. (2013) used MVPA in the context of a target detection task and found that distributed patterns of BOLD activity in PrC honored differences between faces, scenes, words, and sounds, with face representations being significantly different from all other types of

stimulus categories examined (see also Diana et al., 2008; Huffman and Stark, 2014). Our study extends this prior research by showing category specificity in memory signals at retrieval that are directly related to participants' recognition-memory responses. That such spatially distributed signals in PrC show category-specificity in their organization is also suggested by the lack of successful cross-classification across different stimulus categories that we noted in our initial report on this study (Martin et al., 2013). Specifically, these earlier analyses revealed that the linear classifier that had been trained successfully for classification of the familiarity of faces was unsuccessful when tested for classification of the familiarity of chairs even though the familiarity of chairs could be decoded based on right PrC activity when a classifier was directly trained with stimuli from this category. Further, our earlier analyses (unreported results) also showed that when trials with stimuli from both categories were combined during training and test, we were unable to classify corresponding recognition-memory decisions.

Our findings on category-specific recognition-memory signals in PrC are of direct relevance for a representational framework of MTL organization as proposed in various theoretical accounts (Murray and Bussey, 1999; Bussey and Saksida, 2007; Graham et al. 2010; Cowell et al. 2010). Within this framework, PrC is considered to be the apex of the ventral visual pathway in that it carries information about complex objects at the highest level of specificity. From this perspective, the link between PrC functioning and familiarity arises due to the nature of stimuli used, and due to the fidelity of stimulus representations required to succeed with item-based memory discriminations in typical recognition-memory tasks. However, PrC should also play a role in nonmnemonic task contexts when they involve the representation of objects at a comparable level of specificity. Indeed, to the extent that our functional localizer task did not have explicit mnemonic demands, the current findings point to correspondence between voxel patterns that code the perceived familiarity of faces and the representation of faces more generally. Past fMRI research provides additional support for this interpretation (see Collins and Olson, 2014, for review), including several studies that revealed an involvement of PrC in the perceptual discrimination of faces under conditions that required fine-grained distinctions between multiple exemplars, as well as in perceptual learning of faces (Lee et al., 2008; O'Neil et al. 2009; Barense et al., 2010; Mundy et al., 2012, 2013; O'Neil et al., 2013). For example, O'Neil et al. (2009) reported that activity in right PrC is related not only to the accuracy of forced-choice recognition-memory judgments, but also to the accuracy of perceptual oddity judgments for faces.

The fact that we were able to train a linear classifier to decode the perceived familiarity of faces on a set of trials with old and new faces, and then perform accurately in classification on other non-trained test trials implies that activity patterns in PrC generalize, at least in part, across different exemplars of familiar and different exemplars of novel faces, respectively. This regularity may, at first glance, appear difficult to reconcile with the representational account of PrC functioning, which, as discussed, emphasizes a role

of this structure in the representation of objects and faces at a high level of specificity (Bussey and Saksida, 2007; Graham et al. 2010; Cowell et al. 2010). We note, however, that the generalization we observed does not imply that the specificity required for recognition of unique face exemplars could not be retained in PrC activity. Overlap in PrC response patterns for different face exemplars with the same memory status may have been large but incomplete. Indeed, previous fMRI studies that required discrimination between a much smaller set of stimuli, each presented many times, have shown that it is possible to decode the identity of specific face exemplars from activity patterns in the anterior temporal lobe, including from a region in the anterior collateral sulcus that is part of PrC (Kriegeskorte et al., 2007; Nestor et al., 2011; Anzellotti et al., 2013; Anzellotti and Caramazza, 2014; Ramon et al., 2015). Future research can build on the present findings and track how the patterns of activity that we report emerge as specific faces change from being perceived as novel to being familiar.

A final issue that warrants consideration is whether the coding properties we revealed in the current study generalize to situations in which targets and distractors (i.e., old vs. new items) have distinct semantic meaning as in classic word-list learning experiments. Notably, the majority of previous fMRI studies that have linked item recognition to response reductions in contiguous PrC clusters at retrieval have used stimuli with semantic meaning that had discriminative value for the task at hand (see Diana et al., 2007; Kim, 2013, for reviews). By contrast, the faces used in this study were completely novel to all participants and were not directly associated with any semantic knowledge that would easily allow for their discrimination (i.e., they were not faces of famous individuals). The relatively consistent decrease in focal PrC signal revealed in prior studies on recognition memory with words may be a reflection of the influence of semantic knowledge on discrimination. Such an interpretation would be in line with the proposal that the PrC response reductions reported in these studies may reflect conceptual fluency (Voss et al., 2009; Wang et al., 2010; Dew and Cabeza, 2013; Heusser et al., 2013; Wang et al., 2014). Against this background, it is possible that the extent to which item-recognition signals in PrC are reflected in distributed response patterns or in more focal clusters (with a homogeneous direction of response to recent exposure) relates to whether recognition decisions are based primarily on perceptual or conceptual stimulus properties, respectively. This notion could be evaluated in future fMRI research with a direct comparison of the coding properties of recognition-memory signals for stimuli that vary with respect to the discriminative value of perceptual and conceptual item features.

REFERENCES

- Aggleton JP, Brown MW, Albasser MM. 2012. Contrasting brain activity patterns for item recognition memory and associative recognition memory: Insights from immediate-early gene functional imaging. *Neuropsychologia* 50:3141–3155.
- Anzellotti S, Caramazza A. 2014. The neural mechanisms for the recognition of face identity in humans. *FrontPsychol* 5:672.

- Anzellotti S, Fairhall SL, Caramazza A. 2013. Decoding representations of face identity that are tolerant to rotation. *Cerebr Cortex* 24:1988–1995.
- Barense MD, Henson RN, Lee AC, Graham KS. 2010. Medial temporal lobe activity during complex discrimination of faces, objects, and scenes: Effects of viewpoint. *Hippocampus* 20:389–401.
- Behrmann M, Plaut DC. 2013. Distributed circuits, not circumscribed centers, mediate visual recognition. *Trends Cogn Sci* 17:210–219.
- Brown MW, Aggleton JP. 2001. Recognition memory: What are the roles of the perirhinal cortex and hippocampus? *Nat Rev Neurosci* 2:51–61.
- Brown MW, Wilson FAW, Riches IP. 1987. Neuronal evidence that inferomedial temporal cortex is more important than hippocampus in certain processes underlying recognition memory. *Brain Res* 409:158–162.
- Brown MW, Barker GRI, Aggleton JP, Warburton EC. 2012. What pharmacological interventions indicate concerning the role of the perirhinal cortex in recognition memory. *Neuropsychologia* 50:3122–3140.
- Burke SN, Maurer AP, Nematollahi S, Uprety A, Wallace JL, Barnes CA. 2014. Advanced age dissociates dual functions of the perirhinal cortex. *J Neurosci* 34:467–480.
- Bussey TJ, Saksida LM. 2007. Memory, perception, and the ventral visual-perirhinal-hippocampal stream: Thinking outside of the boxes. *Hippocampus* 17:898–908.
- Cate AD, Goodale MA, Köhler S. 2011. The role of apparent size in building-and object-specific regions of ventral visual cortex. *Brain Res* 1388:109–122.
- Collins JA, Olson IR. 2014. Beyond the FFA: The role of the ventral anterior temporal lobes in face processing. *Neuropsychologia* 61:65–79.
- Coutanche MN. 2013. Distinguishing multi-voxel patterns and mean activation: Why, how, and what does it tell us? *Cogn Affect Behav Neurosci* 13:667–673.
- Cowell RA, Bussey TJ, Saksida LM. 2006. Why does brain damage impair memory? A connectionist model of object recognition memory in perirhinal cortex. *J Neurosci* 26:12186–12197.
- Cowell RA, Bussey TJ, Saksida LM. 2010. Components of recognition memory: Dissociable cognitive processes or just differences in representational complexity? *Hippocampus* 20:1245–1262.
- Danckert SL, Gati JS, Menon RS, Köhler S. 2007. Perirhinal and hippocampal contributions to visual recognition memory can be distinguished from those of occipito-temporal structures based on conscious awareness of prior occurrence. *Hippocampus* 17:1081–1092.
- Daselaar SM, Fleck MS, Cabeza R. 2006a. Triple dissociation in the medial temporal lobes: Recollection, familiarity, and novelty. *J Neurophysiol* 96:1902–1911.
- Daselaar SM, Fleck MS, Prince SE, Cabeza R. 2006b. The medial temporal lobe distinguishes old from new independently of consciousness. *J Neurosci* 26:5835–5839.
- Davis T, LaRocque KF, Mumford JA, Norman KA, Wagner AD, Poldrack RA. 2014. What do differences between multi-voxel and univariate analysis mean? How subject-, voxel-, and trial-level variance impact fMRI analysis. *NeuroImage* 97:271–283.
- de Gardelle V, Waszczuk M, Egner T, Summerfield C. 2013. Concurrent repetition enhancement and suppression responses in extrastriate visual cortex. *Cerebr Cortex* 23:2235–2244.
- Desimone R. 1996. Neural mechanisms for visual memory and their role in attention. *Proc Natl Acad Sci USA* 93:13494–13499.
- Dew IT, Cabeza R. 2013. A broader view of perirhinal function: From recognition memory to fluency-based decisions. *J Neurosci* 33:14466–14474.
- Diana RA, Yonelinas AP, Ranganath C. 2007. Imaging recollection and familiarity in the medial temporal lobe: A three-component model. *Trends Cogn Sci* 11:379–386.
- Diana RA, Yonelinas AP, Ranganath C. 2008. High-resolution multi-voxel pattern analysis of category selectivity in the medial temporal lobes. *Hippocampus* 18:536–541.
- Diedrichsen J, Wiestler T, Ejaz N. 2013. A multivariate method to determine the dimensionality of neural representation from population activity. *Neuroimage* 76:225–235.
- Dobbins IG, Han S. 2006. Cue-versus probe-dependent prefrontal cortex activity during contextual remembering. *J Cogn Neurosci* 18:1439–1452.
- Eacott MJ, Gaffan D, Murray EA. 1994. Preserved recognition memory for small sets, and impaired stimulus identification for large sets, following rhinal cortex ablations in monkeys. *Eur J Neurosci* 6:1466–1478.
- Eichenbaum H, Yonelinas AP, Ranganath C. 2007. The medial temporal lobe and recognition memory. *Annu Rev Neurosci* 30:123–152.
- Friston KJ. 1998. Imaging neuroscience: Principles or maps? *Proc Natl Acad Sci USA* 95:796–802.
- Ganel T, Gonzalez CL, Valyear KF, Culham JC, Goodale MA, Köhler S. 2006. The relationship between fMRI adaptation and repetition priming. *Neuroimage* 32:1432–1440.
- Gonsalves BD, Kahn I, Curran T, Norman KA, Wagner AD. 2005. Memory strength and repetition suppression: Multimodal imaging of medial temporal cortical contributions to recognition. *Neuron* 47:751–761.
- Gotts SJ, Chow CC, Martin A. 2012. Repetition priming and repetition suppression: A case for enhanced efficiency through neural synchronization. *Cogn Neurosci* 3:227–237.
- Graham KS, Barense MD, Lee AC. 2010. Going beyond LTM in the MTL: A synthesis of neuropsychological and neuroimaging findings on the role of the medial temporal lobe in memory and perception. *Neuropsychologia* 48:831–853.
- Grill-Spector K, Weiner KS. 2014. The functional architecture of the ventral temporal cortex and its role in categorization. *Nat Rev Neurosci* 15:536–548.
- Grill-Spector K, Henson RNA, Martin A. 2006. Repetition and the brain: Neural models of stimulus-specific effects. *Trends Cogn Sci* 10:14–23.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293:2425–2430.
- Haxby JV, Connolly AC, Guntupalli JS. 2014. Decoding neural representational spaces using multivariate pattern analysis. *Annu Rev Neurosci* 37:435–456.
- Henson RNA, Rugg MD. 2003. Neural response suppression, haemodynamic repetition effects, and behavioural priming. *Neuropsychologia* 41:263–270.
- Henson RNA, Gagnepain P. 2010. Predictive, interactive multiple memory systems. *Hippocampus* 20:1315–1326.
- Henson RNA, Cansino S, Herron JE, Robb WGK, Rugg MD. 2003. A familiarity signal in human anterior medial temporal cortex? *Hippocampus* 13:301–304.
- Heusser AC, Awipi T, Davachi L. 2013. The ups and downs of repetition: Modulation of the perirhinal cortex by conceptual repetition predicts priming and long-term memory. *Neuropsychologia* 51:2333–2343.
- Hölscher C, Rolls ET, Xiang J. 2003. Perirhinal cortex neuronal activity related to long-term familiarity memory in the macaque. *Eur J Neurosci* 18:2037–2046.
- Huffman DJ, Stark CE. 2014. Multivariate pattern analysis of the human medial temporal lobe revealed representationally categorical cortex and representationally agnostic hippocampus. *Hippocampus* 24:1394–1403.
- Hung CP, Kreiman G, Poggio T, DiCarlo JJ. 2005. Fast readout of object identity from macaque inferior temporal cortex. *Science* 310:863–866.

- Kafkas A, Montaldi D. 2012. Familiarity and recollection produce distinct eye movement, pupil and medial temporal lobe responses when memory strength is matched. *Neuropsychologia* 50:3080–3093.
- Kim H. 2013. Differential neural activity in the recognition of old versus new events: An activation likelihood estimation meta-analysis. *Hum Brain Mapp* 34:814–836.
- Kriegeskorte N, Formisano E, Sorger B, Goebel R. 2007. Individual faces elicit distinct response patterns in human anterior temporal cortex. *Proc Natl Acad Sci USA* 104:20600–20605.
- Kriegeskorte N, Mur M, Ruff DA, Kiani R, Bodurka J, Esteky H, Tanaka K, Bandettini PA. 2008. Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* 60:1126–1141.
- Lee ACH, Buckley MJ, Gaffan D, Emery T, Hodges JR, Graham KS. 2006. Differentiating the roles of the hippocampus and perirhinal cortex in processes beyond long-term declarative memory: A double dissociation in dementia. *J Neurosci* 26:5198–5203.
- Lee ACH, Scabill VL, Graham KS. 2008. Activating the medial temporal lobe during oddity judgment for faces and scenes. *Cerebr Cortex* 18:683–696.
- Lee JH, Durand R, Gradinaru V, Zhang F, Goshen I, Kim DS, Fenno LE, Ramakrishnan C, Deisseroth K. 2010. Global and local fMRI signals driven by neurons defined optogenetically by type and wiring. *Nature* 465:788–792.
- Liang JC, Wagner AD, Preston AR. 2013. Content representation in the human medial temporal lobe. *Cerebr Cortex* 23:80–96.
- Logothetis NK. 2008. What we can do and what we cannot do with fMRI. *Nature* 453:869–878.
- Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412:150–157.
- Martin CB, Mirsattari SM, Pruessner JC, Pietrantonio S, Burneo JG, Hayman-Abello B, Köhler S. 2012. Déjà vu in unilateral temporal-lobe epilepsy is associated with selective familiarity impairments on experimental tasks of recognition memory. *Neuropsychologia* 50:2981–2991.
- Martin CB, McLean DA, O'Neil EB, Köhler S. 2013. Distinct familiarity-based response patterns for faces and buildings in perirhinal and parahippocampal cortex. *J Neurosci* 33:10915–10923.
- Meunier M, Bachevalier J, Mishkin M, Murray EA. 1993. Effects on visual recognition of combined and separate ablations of the entorhinal and perirhinal cortex in rhesus monkeys. *J Neurosci* 13:5418–5432.
- Misaki M, Kim Y, Bandettini PA, Kriegeskorte N. 2010. Comparison of multivariate classifiers and response normalizations for pattern-information fMRI. *Neuroimage* 53:103–118.
- Montaldi D, Mayes AR. 2010. The role of recollection and familiarity in the functional differentiation of the medial temporal lobes. *Hippocampus* 20:1291–1314.
- Montaldi D, Spencer TJ, Roberts N, Mayes AR. 2006. The neural system that mediates familiarity memory. *Hippocampus* 16:504–520.
- Mundy ME, Downing PE, Graham KS. 2012. Extrastriate cortex and medial temporal lobe regions respond differentially to visual feature overlap within preferred stimulus category. *Neuropsychologia* 50:3053–3061.
- Mundy ME, Downing PE, Dwyer DM, Honey RC, Graham KS. 2013. A critical role for the hippocampus and perirhinal cortex in perceptual learning of scenes and faces: Complementary findings from amnesia and fMRI. *J Neurosci* 33:10490–10502.
- Murray EA, Bussey TJ. 1999. Perceptual–mnemonic functions of the perirhinal cortex. *Trends Cogn Sci* 3:142–151.
- Murray EA, Bussey TJ, Saksida LM. 2007. Visual perception and memory: A new view of medial temporal lobe function in primates and rodents. *Annu Rev Neurosci* 30:99–122.
- Nasr S, Tootell RB. 2012. Role of fusiform and anterior temporal cortical areas in facial recognition. *Neuroimage* 63:1743–1753.
- Nestor A, Plaut DC, Behrmann M. 2011. Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. *Proc Natl Acad Sci USA* 108:9998–10003.
- Norman KA, O'Reilly RC. 2003. Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. *Psychol Rev* 110:611–646.
- Norman KA, Polyn SM, Detre GJ, Haxby JV. 2006. Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10:424–430.
- O'Neil EB, Cate AD, Köhler S. 2009. Perirhinal cortex contributes to accuracy in recognition memory and perceptual discriminations. *J Neurosci* 29:8329–8334.
- O'Neil EB, Barkley VA, Köhler S. 2013. Representational demands modulate involvement of perirhinal cortex in face processing. *Hippocampus* 23:592–605.
- O'Neil EB, Hutchison RM, McLean DA, Köhler S. 2014. Resting-state fMRI reveals functional connectivity between face-selective perirhinal cortex and the fusiform face area related to face inversion. *Neuroimage* 92:349–355.
- Op de Beeck HPO, Haushofer J, Kanwisher NG. 2008. Interpreting fMRI data: Maps, modules and dimensions. *Nat Rev Neurosci* 9:123–135.
- Pruessner JC, Li LM, Serles W, Pruessner M, Collins DL, Kabani N, Lupien S, Evans AC. 2000. Volumetry of hippocampus and amygdala with high-resolution MRI and three-dimensional analysis software: Minimizing the discrepancies between laboratories. *Cerebr Cortex* 10:433–442.
- Pruessner JC, Köhler S, Crane J, Pruessner M, Lord C, Byrne A, Kabani N, Collins DL, Evans AC. 2002. Volumetry of temporopolar, perirhinal, entorhinal and parahippocampal cortex from high-resolution MR images: Considering the variability of the collateral sulcus. *Cerebr Cortex* 12:1342–1353.
- Quamme JR, Weiss DJ, Norman KA. 2010. Listening for recollection: A multi-voxel pattern analysis of recognition memory retrieval strategies. *Front Hum Neurosci* 4:61.
- Rajimehr R, Young JC, Tootell RB. 2009. An anterior temporal face patch in human cortex, predicted by macaque maps. *Proc Natl Acad Sci USA* 106:1995–2000.
- Ramon M, Vizioli L, Liu-Shuang J, Rossion B. 2015. Neural microgenesis of personally familiar face recognition. *Proc Natl Acad Sci USA* 112:E4835–E4844.
- Ranganath C. 2010. A unified framework for the functional organization of the medial temporal lobes and the phenomenology of episodic memory. *Hippocampus* 20:1263–1290.
- Ringo JL. 1996. Stimulus specific adaptation in inferior temporal and medial temporal cortex of the monkey. *Behav Brain Res* 76:191–197.
- Rissman J, Wagner AD. 2012. Distributed representations in memory: insights from functional brain imaging. *Annu Rev Psychol* 63:101–128.
- Rossion B, Hanseeuw B, Dricot L. 2012. Defining face perception areas in the human brain: a large-scale factorial fMRI face localizer analysis. *Brain Cogn* 79:138–157.
- Segaert K, Weber K, de Lange FP, Petersson KM, Hagoort P. 2013. The suppression of repetition enhancement: a review of fMRI studies. *Neuropsychologia* 51:59–66.
- Sobotka S, Ringo JL. 1996. Mnemonic responses of single units recorded from monkey inferotemporal cortex, accessed via trans-commissural versus direct pathways: a dissociation between unit activity and behavior. *J Neurosci* 16:4222–4230.
- Staresina BP, Duncan KD, Davachi L. 2011. Perirhinal and parahippocampal cortices differentially contribute to later recollection of object- and scene-related event details. *J Neurosci* 31:8739–8747.
- Squire LR, Zola-Morgan M, Zola-Morgan J, Clark RE. 2007. Recognition memory and the medial temporal lobe: A new perspective. *Nat Rev Neurosci* 8:872–883.

- Thome A, Erickson CA, Lipa P, Barnes CA. 2012. Differential effects of experience on tuning properties of macaque MTL neurons in a passive viewing task. *Hippocampus* 22:2000–2011.
- Tong F, Pratte MS. 2012. Decoding patterns of human brain activity. *Annu Rev Psychol* 63:483–509.
- Voss JL, Hauner KK, Paller KA. 2009. Establishing a relationship between activity reduction in human perirhinal cortex and priming. *Hippocampus* 19:773–778.
- Wang WC, Lazzara MM, Ranganath C, Knight RT, Yonelinas AP. 2010. The medial temporal lobe supports conceptual implicit memory. *Neuron* 68:835–842.
- Wang WC, Ranganath C, Yonelinas AP. 2014. Activity reductions in perirhinal cortex predict conceptual priming and familiarity-based recognition. *Neuropsychologia* 52:19–26.
- Wixted JT, Mickes L, Squire LR. 2010. Measuring recollection and familiarity in the medial temporal lobe. *Hippocampus* 20:1195–1205.
- Xiang JZ, Brown MW. 1998. Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology* 37:657–676.
- Yassa MA, Stark CE. 2008. Multiple signals of recognition memory in the medial temporal lobe. *Hippocampus* 18:945–954.
- Yonelinas AP. 1999. The contribution of recollection and familiarity to recognition and source-memory judgments: A formal dual-process model and an analysis of receiver operating characteristics. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 25:1415–1434.
- Yonelinas AP, Aly M, Wang WC, Koen JD. 2010. Recollection and familiarity: Examining controversial assumptions and new directions. *Hippocampus* 20:1178–1194.
- Zhu XO, Brown MW, Aggleton JP. 1995. Neuronal signaling of information important to visual recognition memory in rat rhinal and neighbouring cortices. *Eur J Neurosci* 7:753–765.