

# Recognition Alters the Spatial Pattern of fMRI Activation in Early Retinotopic Cortex

P.-J. Hsieh, E. Vul, and N. Kanwisher

Department of Brain and Cognitive Sciences, McGovern Institute, Massachusetts Institute of Technology, Cambridge, Massachusetts

Submitted 4 September 2009; accepted in final form 11 January 2010

**Hsieh P.-J., Vul E., Kanwisher N.** Recognition alters the spatial pattern of fMRI activation in early retinotopic cortex. *J Neurophysiol* 103: 1501–1507, 2010. First published January 13, 2010; doi:10.1152/jn.00812.2009. Early retinotopic cortex has traditionally been viewed as containing a veridical representation of the low-level properties of the image, not imbued by high-level interpretation and meaning. Yet several recent results indicate that neural representations in early retinotopic cortex reflect not just the sensory properties of the image, but also the perceived size and brightness of image regions. Here we used functional magnetic resonance imaging pattern analyses to ask whether the representation of an object in early retinotopic cortex changes when the object is recognized compared with when the same stimulus is presented but not recognized. Our data confirmed this hypothesis: the pattern of response in early retinotopic visual cortex to a two-tone “Mooney” image of an object was more similar to the response to the full grayscale photo version of the same image when observers knew what the two-tone image represented than when they did not. Further, in a second experiment, high-level interpretations actually overrode bottom-up stimulus information, such that the pattern of response in early retinotopic cortex to an identified two-tone image was more similar to the response to the photographic version of that stimulus than it was to the response to the identical two-tone image when it was not identified. Our findings are consistent with prior results indicating that perceived size and brightness affect representations in early retinotopic visual cortex and, further, show that even higher-level information—knowledge of object identity—also affects the representation of an object in early retinotopic cortex.

## INTRODUCTION

Early retinotopic cortex (i.e., V1/V2/V3) has traditionally been viewed as an approximately veridical representation of the low-level properties of the image, which provides the input to higher cortical areas where meaning is assigned to the image. But how distinct are data and theory in visual cortex? Recent evidence has shown that activity in early retinotopic cortex can reflect midlevel visual computations, such as contour completion (Stanley and Rubin 2003), figure-ground discrimination (Heinen et al. 2005; Huang and Paradiso 2008; Hupé et al. 1998), perceived size (Murray et al. 2006) and lightness constancy (Boyaci et al. 2007). Further, at least some of the responses in early retinotopic cortex are likely affected by feedback from higher areas, including attentional modulations (Datta and DeYoe 2009; Fischer and Whitney 2009; Ress et al. 2000), stimulus reward history information (Serences 2008), and lower responses when a shape is perceived compared with when it is not (Fang et al. 2008a; Murray et al. 2002). Here we ask whether retinotopic cortex contains infor-

mation about the perceived identity of an object. We uncovered the perceived identity of the object from its low-level visual properties by comparing the response to a Mooney image before and after it is disambiguated.

Specifically, we scanned subjects while they viewed Mooney images they could not interpret, then grayscale versions of those same images, then the original Mooney images, which subjects could now understand. We found that the spatial pattern of activation in early retinotopic cortex of a Mooney image more closely resembles the activation pattern from the unambiguous grayscale version of that same image when the Mooney image is understood than when it is not. This finding goes beyond prior findings indicating that the representation of a stimulus in early retinotopic cortex reflects higher-level information such as perceived size (Fang et al. 2008a; Murray et al. 2002), brightness (Boyaci et al. 2007), and grouping (Fang et al. 2008b; Murray et al. 2002, 2004) by showing that the activation pattern of a stimulus in early retinotopic cortex also changes when it is recognized compared with when it is not. Further, our use of pattern analysis enables us to show that 1) the specific way the activation pattern (i.e., representation) in early retinotopic cortex changes is to become more like the unambiguous stimulus and 2) these higher-level effects on pattern information are dissociable from the mean changes that have been reported previously.

## METHODS

### Participants

Seven subjects (age range: 20–30 yr) were run in the blocked design experiment (*experiment 1*) and eight subjects (age range: 20–30 yr) were run in the event-related design experiment (*experiment 2*). All subjects had normal or corrected-to-normal vision. All gave written consent within a protocol passed by the Massachusetts Institute of Technology committee on the Use of Humans as Experimental Subjects. Subjects were paid \$60.00 per session.

### Experimental procedures

In *experiment 1* (Fig. 1A), each subject participated in 12 runs. In the first 4 runs (phase1: the “Mooney1” condition), subjects were presented with two ambiguous two-tone images containing unidentifiable objects ( $7 \times 7$  in visual degrees). In the following 4 runs (phase2: “Photo”), subjects were presented with grayscale versions of the same stimuli in which the object could be clearly identified. In the last 4 runs (phase3: “Mooney2”), the original two-tone images were presented again, but now the pictured object could be easily recognized due to the experience of the grayscale images. In all, there were eight blocks of stimulus presentation in each run. During each block (20 s), the same image (a fish or a camel) was presented at 1 Hz at the center of the screen (remaining on for 500 ms, followed by a blank interstimulus interval of 500 ms). The two kinds of block (fish-block

Address for reprint requests and other correspondence: P.-J. Hsieh, Department of Brain and Cognitive Sciences, McGovern Institute, Massachusetts Institute of Technology, Cambridge, MA 02139 (E-mail: pjh@mit.edu).

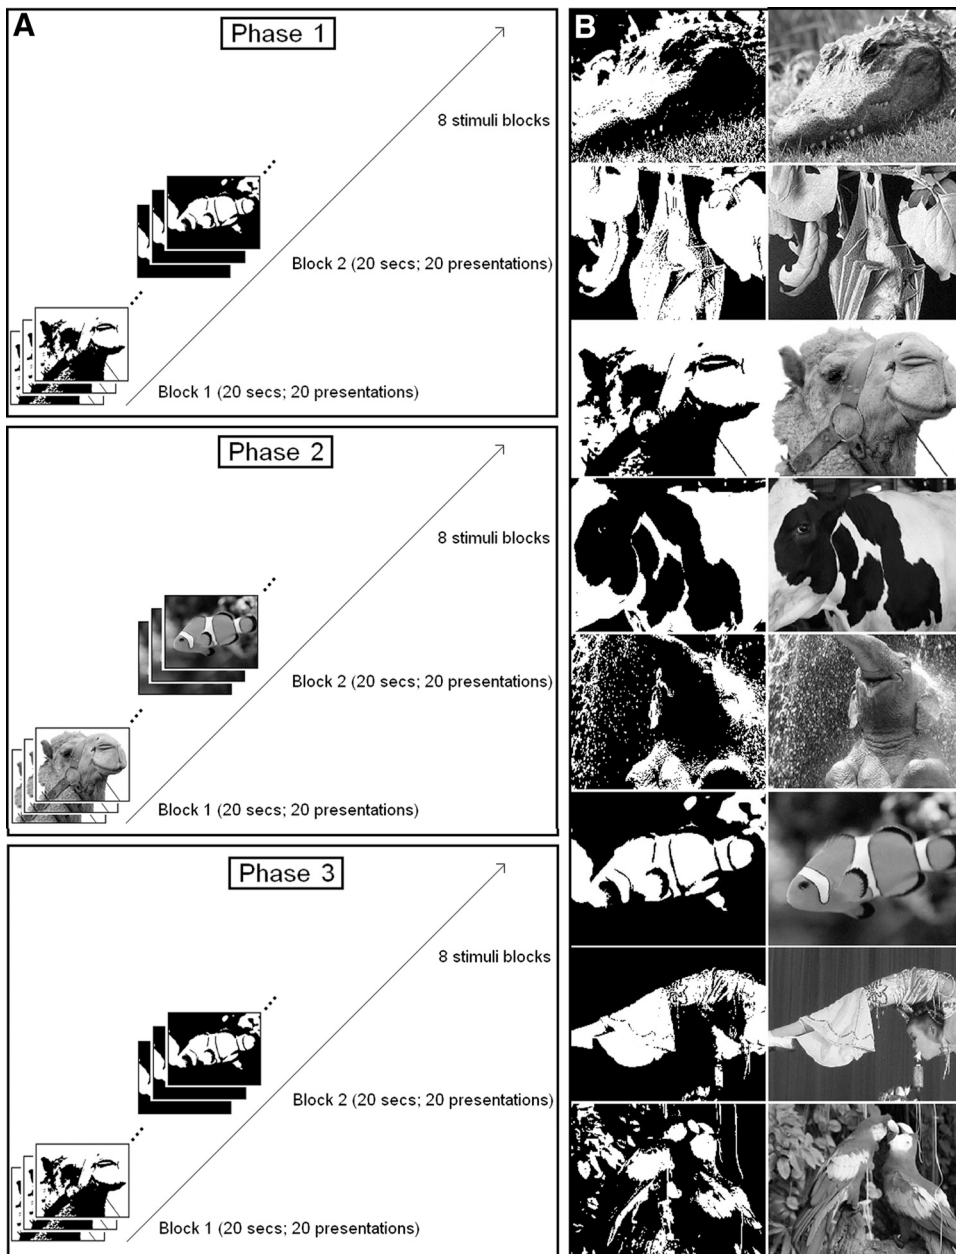


FIG. 1. Experimental procedure (blocked design). A: during 3 phases of the experiment subjects viewed 1) 2 different 2-tone images that they could not identify ("Mooney1"); 2) the easily identifiable grayscale versions of the same images ("Photo"); 3) the original 2-tone images that could now be easily recognized because of experience with the corresponding photograph ("Mooney2"). B: in *experiment 2*, 2 images were randomly selected for each subject from the set of 8 images shown here.

and camel-block) were interleaved with blank blocks (20 s each); the order of blocks was counterbalanced across runs. Each picture had two possible contrasts (high or low). For each stimulus presentation, participants were required to indicate via a two-button response box whether the stimulus had a relatively high or low contrast.

In *experiment 2*, all the procedures were identical to those in *experiment 1*, except there were 60 events total of stimulus presentation in each run (30 for each stimulus condition). During each stimulus presentation, one of the two images was presented at the center of the screen for 180 ms. The order of trials was optimized using Optseq2, an optimal sequencing program (NMR Center, Massachusetts General Hospital, Boston, MA). Each run lasted 5 min 15 s. More than one third of the scanning time (3/7) consisted of null events that were randomly inserted between trials. For each subject, two images were randomly selected from a set of eight images (Fig. 1B). Before the experiment, subjects were shown briefly with the two randomly selected pictures (remembered by subjects as pictures 1 and 2) and verbally confirmed that they could not recognize the images. After subjects finished phase1 and before they started phase2, we

presented the photo version of the images during the break to ensure they could recognize the objects in the photo images. After subjects finished phase2 and before they started phase3, we presented the photo version of the images and the Mooney images side by side during the second break to ensure they could recognize that the objects in the Mooney images were the same as those in the Photo images. Subjects were also asked after phase3 if they could identify the stimuli during phase3. For each stimulus presentation during the experiment, participants were required to indicate which one of the two images (picture 1 or picture 2) was presented via a two-button response box.

#### Identification of regions of interest

The localizer scans were run as previously described (Kourtzi and Kanwisher 2001). Functional localization of the regions of interest (ROIs) was based on independent runs (minimum = 1; maximum = 3) of four 20-s blocks of grayscale faces, scenes, common objects, and scrambled objects (four blocks per category per run). The critical ROI




FIG. 2. Regions of interests (ROIs). *A*: a statistical map for lateral occipital complex (LOC, defined as the region that responded more strongly to images of intact objects than to images of scrambled objects;  $P < 10^{-4}$ ) is shown for one representative subject. We selected the 2 regions at the lateral occipital lobe (circled) as LOC. *B*: inflated cortical surfaces of a brain with dark gray regions representing sulci and light gray regions representing gyri. Foveal confluence (FC) was defined as a small region at the posterior end of the calcarine sulcus (green), functionally constrained to the region (yellow) that responded more strongly to images of intact/scramble objects than to fixation ( $P < 10^{-25}$ ). Another annulus ROI surrounding the FC was also defined (red).

identified in the independent localizer scan was the LOC (defined as the region that responded more strongly to images of intact objects than to images of scrambled objects; Malach et al. 1995). The foveal confluence (Dougherty et al. 2003) was defined as a small region at the posterior end of the calcarine sulcus, functionally constrained to the regions that responded more strongly to any images of (intact or scrambled) than to fixation (Fig. 2). Specifically, we overlaid each subject's functional contrast map (intact objects > scrambled objects) on top of his/her inflated brain and then selected voxels located 1) around the posterior tip of the calcarine sulcus and 2) within the functional contrast map. Since our stimuli extended about  $3.5^\circ$  from the fovea and most of the features in the image were located in the center, we decided to be conservative and restricted the size of the foveal confluence so that it was smaller than the minimum size of foveal confluence reported in Dougherty et al. (2003). An additional annulus region surrounding the foveal confluence was identified as a control ROI. Since the annulus ROI was chosen to be a control region, the size of this ROI was selected to be larger than the foveal confluence (to gain power for the correlation analysis). Average sizes of the foveal confluence ROI in *experiments* 1 and 2 were  $161 \pm 19$  (mean  $\pm$  SE; range: 103–257) and  $153 \pm 15$  (range: 103–228) voxels. Pictures for individual ROIs are provided in Supplemental Fig. S1.<sup>1</sup>

### Data analysis

Scanning was done on the 3T Siemens Trio scanner in the McGovern Institute at MIT in Cambridge (Athina A. Martinos Imaging Center). Functional magnetic resonance imaging (fMRI) runs were acquired using a gradient-echo echo-planar sequence (repetition time = 3 s, time to echo = 40 ms,  $1.5 \times 1.5 \times 1.5$  mm + 10% spacing). Forty slices were collected with a 32-channel head coil. Slices were oriented roughly perpendicular to the calcarine sulcus and covered not only most of the occipital and posterior temporal lobes, but also some of the parietal lobes.

fMRI data analysis was conducted using freesurfer (available at <http://surfer.nmr.mgh.harvard.edu/>). The processing steps for both the localizer and the experimental runs included motion correction and linear trend removal. The processing for the localizer also

included spatial smoothing with a 6-mm kernel. A gamma function with delta ( $\delta$ ) = 2.25 and tau ( $\tau$ ) = 1.25 was used to estimate the hemodynamic response for each condition in the localizer scans.

Correlation analysis was conducted on the beta weight for each condition in each voxel with a standard multivariate pattern analysis method (Haxby et al. 2001). The mean response in each voxel across all conditions was subtracted from the response to each individual condition for each half of the data before calculating the correlations. Spatial patterns were extracted from each set of data (Mooney1, Photo, and Mooney2) separately for each combination of ROI and stimulus type (e.g., fish vs. camel). Within each ROI we then computed the correlation between the spatial patterns of Mooney1 and Photo runs from the same stimulus category (e.g., between Mooney1–fish and Photo–fish, and between Mooney1–camel and Photo–camel). The same correlation was computed between the spatial patterns of Mooney2 and Photo runs and between the spatial patterns of Mooney2 and Mooney1 runs. These correlations were computed for each subject and then averaged across subjects by conditions. Note that our method is equivalent to that used by Haxby et al. (2001), in which a given ROI is deemed to contain information about a given stimulus discrimination (stimulus A vs. stimulus B) if the pattern of response across voxels in that region is more similar for two response patterns when they were produced by the same stimulus than when they were produced by two different stimuli: i.e., if  $\text{mean}[r(A1, A2), r(B1, B2)] > \text{mean}[r(A1, B2), r(A2, B1)]$ . In our method, we directly computed the correlation between the spatial patterns of Mooney and Photo runs from the same stimulus category (e.g., between Mooney–fish and Photo–fish and between Mooney–camel and Photo–camel). Although, conceptually, saying that an ROI contains discriminative information about stimulus A versus stimulus B (as in Haxby et al. 2001) is very different from saying that the activation pattern in this ROI is more similar between phase2 and phase3 than that between phase1 and phase2 as in our method, mathematically they amount to the same thing because of the normalization procedure (activation values were normalized to a mean of zero in each voxel across categories by subtracting the mean response across all categories) and the fact that there were only two categories.

<sup>1</sup> The online version of this article contains supplemental data.




FIG. 3. Blocked design results. *A*: the correlation for the 3 comparisons (Mooney1–Photo, Mooney2–Photo, and Mooney2–Mooney1) in LOC. The correlation for Mooney2–Photo is significantly greater than that for Mooney1–Photo ( $P = 0.010$ ). *B*: the correlation for the 3 comparisons in the foveal confluence ROI. The correlation for Mooney2–Photo is significantly greater than that for Mooney1–Photo ( $P = 0.048$ ). Error bars indicate SEs.

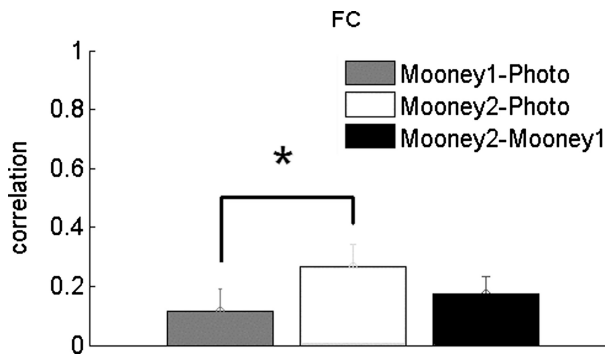


FIG. 4. Event-related results. The correlation for the 3 comparisons (Mooney1–Photo, Mooney2–Photo, and Mooney2–Mooney1) in the foveal confluence ROI. The correlation for Mooney2–Photo is significantly greater than that for Mooney1–Photo ( $P < 0.001$ ). Error bars indicate SEs.

## RESULTS

We tested whether top-down signals alter low-level representations to match high-level interpretations in an fMRI experiment with three phases (Fig. 1A). Subjects viewed 1) two different two-tone images that they could not identify (“Mooney1”); 2) the easily identifiable grayscale photographic versions of the same images, e.g., a camel and a fish (“Photo”); and 3) the original two-tone images that could now be easily recognized because of experience with the corresponding photograph (“Mooney2”). If top-down signals make representations in early retinotopic cortex areas more closely match learned interpretations, we would predict that the fMRI response patterns in early retinotopic cortex will be more similar between the Photo and Mooney2 phases for a given stimulus than those between the Photo and Mooney1 phases for the same stimulus.

To address this question, we used a standard multivariate pattern analysis method (see METHODS; Haxby et al. 2001). This analysis was applied to three (independently defined) ROIs: the lateral occipital complex (LOC); the foveal confluence (FC), an early retinotopic area at the occipital pole where the foveal representations of visual areas V1, V2, and V3 converge; and an annulus ROI surrounding the FC (Fig. 2).

### Pattern information in LOC

First, we considered ventral visual area LOC, where the fMRI response has previously been shown to change after subjects recognized Mooney images (Dolan et al. 1997). Consistent with prior work, we found that the pattern of response across voxels in the LOC reflected a subject’s interpretation of object identity in the first, blocked design experiment (see Fig. 1A). That is, the pattern of response for the camel during the

Photo phase was more similar to the camel in the Mooney2 phase than the camel in the Mooney1 phase [Fig. 3A; paired  $t(6) = 3.69$ ,  $P = 0.010$ ]. This result indicates that the representation of a given Mooney image in LOC is more similar to the representation of the photographic version of the same image after subjects learned what that Mooney image represents than before. Note that previous studies showed only that perceptual learning of faces or objects enhanced activity in ventral visual areas (Deshpande et al. 2009; Dolan et al. 1997; George et al. 1999; James et al. 2002; McKeef and Tong 2006). Here we showed for the first time that the activation pattern of a given Mooney image actually changed to become more similar to the activation pattern of the photographic version of the same image.

### Pattern information in early visual cortex

Our key prediction is that this pattern of correlations should exist in early retinotopic cortex as well: i.e., the high-level interpretation of a stimulus will modify the representation of that stimulus in early retinotopic cortex to make it more closely match what would be expected given that interpretation. Confirming our predictions, we find the same effect in early retinotopic cortex: the Mooney2–Photo correlation was greater for corresponding stimuli than the Mooney1–Photo correlation [Fig. 3B; paired  $t(6) = 2.48$ ,  $P = 0.048$ ]. This result indicates that the representation in early retinotopic cortex for an ambiguous Mooney image is more similar to the early retinotopic representation of a grayscale version of the same image when subjects interpret the Mooney as containing the information in the photo than when they do not. Analysis of the correlations for noncorresponding stimuli (camel to fish) reveals that our results do not simply reflect overall higher responses between stimuli in Photo–Mooney2 versus Photo–Mooney1 (see METHODS). Note that for both the foveal confluence ROI and LOC, the Mooney2–Mooney1 correlation is the highest of all the comparisons. This is not surprising given that the bottom-up information is identical in Mooney2 and Mooney1.

Our finding suggests that the higher-level interpretation of a stimulus apparently alters the representation of that stimulus in early retinotopic cortex to make it match the representation consistent with the higher-level interpretation. However, before this conclusion can be accepted, we must address an alternative account subjects may have fixated different parts of the stimuli (e.g., the eyes) when they could identify the stimulus (i.e., in the Photo and Mooney2 phases), compared with when they could not (the Mooney1 phase), making the patterns of early retinotopic response more similar for the photograph and Mooney2 cases (which are both recognized) than either is to the Mooney1 case (which is not recognized). To test this

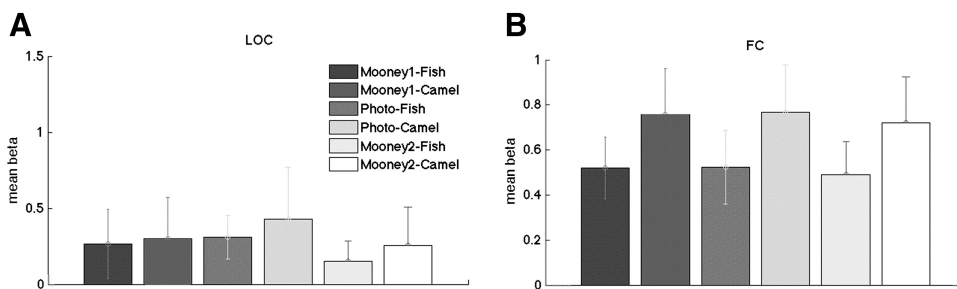


FIG. 5. The mean beta weights of the 6 conditions for experiment 1 in (A) LOC and (B) foveal confluence. Error bars indicate SEs.

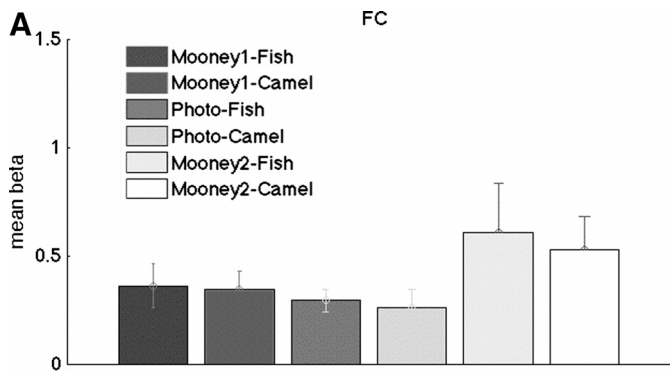


FIG. 6. The mean beta weights of the 6 conditions for *experiment 2* in the foveal confluence ROI. Error bars indicate SEs.

possibility, we conducted a second experiment with an event-related design in which the two different stimuli were interleaved in a random order within each block and each stimulus was displayed for only 180 ms. Thus subjects could not predict which stimulus would appear next and they did not have time to make a saccade while a stimulus was presented. Nonetheless, we replicated the findings from our first experiment: we again found a higher correlation between Photo and Mooney2 than that between Photo and Mooney1 [Fig. 4; paired  $t(7) = 6.77$ ,  $P < 0.001$ ]. Thus our results cannot be explained by differential eye fixation patterns or predictive signals.

With *experiment 2*, we further tested whether the degree to which the early retinotopic cortex representation matches the interpretation (top-down priors) as opposed to the data (bottom-up input) increases as the amount of bottom-up information decreases. We suspected that in the event-related design, when stimuli are presented briefly (180 ms, compared with 500 ms in *experiment 1*) and less stimulus information is available, the early retinotopic representation of Mooney2 would be even more similar to the Photo than to Mooney1. This possibility is consistent with many models of vision positing that V1 reflects a combination of high-level information and low-level properties of an image. For example, one popular proposal about the computational role of top-down signals in the visual system is that these signals are the brain's implementation of hierarchical Bayesian inference; i.e., they are the priors for the low-level visual areas. On this account, if prior information is particularly strong and incoming data are particularly ambiguous, then the patterns of response between Mooney2 and the Photo in early retinotopic cortex might be even more similar than those between Mooney2 and Mooney1.

Indeed, our data are consistent with this interpretation: presenting less bottom-up information (in *experiment 2* compared with *experiment 1*) results in a proportional increase in the Mooney2–Photo correlation relative to the Mooney2–

Mooney1 correlation (chi-square sign test of the interaction between *experiments 1* and *2* and Mooney2–Photo vs. Mooney2–Mooney1,  $P = 0.02$ ). This interaction between the amount of stimulus information and the amount of top-down information supports the hypothesis that the modulation of early retinotopic cortex in this task reflects not just a superposition but an integration of top-down interpretation (priors) with the bottom-up data. However, further work is necessary to test whether bottom-up and top-down inputs to retinotopic cortex interact.

#### No mean response difference in LOC and foveal confluence ROI

For the blocked design experiment (*experiment 1*), the mean responses in both foveal confluence ROI and LOC (mean beta weights averaged within ROI) do not differ between Mooney1 and Mooney2 [paired  $t(6) < 1.23$ ,  $P > 0.26$ ], indicating that the pattern information is independent of the mean response (Fig. 5). For the event-related experiment (*experiment 2*), there is a trend of a higher response in FC for Mooney2 than that for Mooney1 (Fig. 6), but it does not reach significance [paired  $t(7) = 1.86$ ,  $P = 0.11$ ]. This result contrasts with previous studies (McKeef and Tong 2007) that recognized Mooney images induce larger signals in higher visual areas than unrecognized Mooney images. We suspect that this discrepancy arises from the presentation of only two Mooney images in our experiment, which might lead to adaptation effects stronger than those in these prior experiments.

#### Pattern of response did not change in the annulus ROI surrounding the foveal confluence

We checked the anatomical specificity of our effect by analyzing a control ROI, defined as an annulus surrounding the foveal confluence ROI. We found none of the critical effects in this annulus ROI: correlations between Photo and Mooney2 were identical to those between Photo and Mooney1, for both *experiment 1* [paired  $t(6) = 0.07$ ,  $P = 0.94$ , Fig. 7A] and *experiment 2* [paired  $t(7) = 0.70$ ,  $P = 0.51$ , Fig. 7B]. The two-way ANOVA (region  $\times$  experiment) revealed a significant region effect [ $F(1,25) = 11.84$ ,  $P = 0.002$ ]. These results suggest that our finding did not result from a spillover from higher-level areas.

#### DISCUSSION

Our results show that the recognition of an object alters its representation in early retinotopic cortex. Specifically, the early retinotopic representation of an ambiguous Mooney stimulus becomes more similar to the representation of an unam-




FIG. 7. Pattern analysis results in the annulus ROI surrounding the foveal confluence. A: the correlation for the 3 comparisons (Mooney1–Photo, Mooney2–Photo, and Mooney2–Mooney1) in *experiment 1*. The correlation for Mooney2–Photo is not significantly greater than that for Mooney1–Photo ( $P = 0.94$ ). B: the correlation for the 3 comparisons in *experiment 2*. The correlation for Mooney2–Photo is not significantly greater than that for Mooney1–Photo ( $P = 0.51$ ). Error bars indicate SEs.

biguous grayscale version of the same stimulus when the Mooney stimulus is recognized compared with when it is not. Our finding goes beyond prior reports of correlations between neural responses and subjective percepts (McKeeff and Tong 2006; Serences and Boynton 2007a,b), which might reflect trial-to-trial fluctuations in neural noise in early retinotopic visual areas propagating in a feedforward fashion to higher areas, producing different subjective percepts. In contrast, in our study we directly manipulated the high-level interpretation of the stimulus and observed its effect in early retinotopic regions, enabling us to infer the causal direction of our effect as reflecting feedback from higher areas (such as LOC) to early retinotopic visual areas (Murray et al. 2002; Williams et al. 2008). Our findings therefore suggest that representations in early retinotopic visual cortex reflect an integration of high-level information about the identity of a viewed object and low-level properties of that object.

Our finding is consistent with other evidence that responses in early visual cortex can reflect influences of mid- and high-level visual processes, including figure-ground segregation (Heinen et al. 2005; Huang and Paradiso 2008; Hupé et al. 1998), contour completion (Stanley and Rubin 2003), identification of object parts (McKeeff and Tong 2006), perceived size (Fang et al. 2008a; Murray et al. 2006), lightness constancy (Boyaci et al. 2007), and grouping (Fang et al. 2008b; Murray et al. 2002). Our finding goes beyond these other findings by showing an effect of object recognition on the representation of a stimulus in early retinotopic cortex. Further, our use of pattern analysis enabled us to show that the specific nature of change in early retinotopic cortex is that the representation becomes more like that of the unambiguous (grayscale) stimulus when the stimulus is recognized (Mooney2) compared with when it is not (Mooney1).

Perhaps the prior finding closest to ours is the report that mean responses in early retinotopic cortex are higher for stimuli that have not been grouped or recognized than for stimuli that have (Fang et al. 2008a; Murray 2002). We did not find this previously reported difference in mean responses in our study, despite finding that the activation pattern in early retinotopic cortex changes after the subjects recognized the stimuli. This discrepancy between our results and the prior studies may be explained by the fact that our subjects were performing an orthogonal task, whereas some of the previously reported changes in early retinotopic cortex attributed to high-level interpretations are dependent on attention (Fang et al. 2008b). In any case, the fact that we found a change in the pattern response without a change in the mean response indicates that the two phenomena are distinct.

Our findings are consistent with several theoretical frameworks. According to one, the visual system implements hierarchical Bayesian inference (Lee and Mumford 2003; Rao 2005), such that each visual area integrates priors (top-down signals) with new data (bottom-up input) to infer the likely causes of the input. According to the framework of perception as Bayesian inference (Knill and Richards 1996), prior knowledge about the structure of the world provides the additional constraints necessary to solve the problem of vision—that is, to infer the causes in the world of particular retinal images (Lee 1995; Mumford 1992). Although numerous proposals have been made about how Bayesian inference could be implemented in the brain (Lee and Mumford 2003), neural evidence

for such computations has proven elusive (Knill and Pouget 2004). Our findings provide evidence for one specific implementation of Bayesian inference in the visual system: a hierarchical system in which each layer integrates prior information from higher levels of abstractions with data from lower layers of abstraction (Lee and Mumford 2003).

Other interpretations of our finding include the possibility that the changes in early retinotopic cortex that arise when an image is understood reflect interpretation-contingent visual attention to specific retinotopic coordinates or mental imagery (Klein et al. 2004; Mohr et al. 2009). These effects might occur after recognition is complete, in which case they would constitute an alternative to the Bayesian account. However, attention/imagery accounts, according to which these mechanisms play a role in recognition itself, need not be alternatives to the Bayesian account. For example, attentional allocation to specific features at particular spatial locations based on high-level information is one method of combining top-down priors and bottom-up information and can be seen as mechanisms by which hierarchical Bayesian inference is accomplished. Indeed, attention to specific spatial coordinates has been postulated as the mechanism driving contour completion (Stanley and Rubin 2003), identification of object parts (McKeeff and Tong 2006), and figure-ground segregation (Hupé et al. 1998). These computations all require constraints in addition to the bottom-up input and can be computationally formalized as hierarchical Bayesian inference. A higher-level interpretation (e.g., of object identity) is a natural source of additional information that can help to segregate figure from ground, fill in surfaces, complete contours, discern object parts, or otherwise sharpen early visual representations: if we learn that an image represents a face, we can infer where to expect contours, parts, figure, and ground. Moreover, the higher-level interpretation is often easier to determine than the details, both for people and Bayesian models (Kersten et al. 2004). We propose that this “blessing of abstraction” allows the visual system to efficiently parse visual scenes by exploiting its hierarchical structure: higher visual areas can quickly infer the global structure of the scene or object and can then provide these interpretations as priors to the lower visual areas, thus disambiguating otherwise ambiguous input data (Bar 2003; Bar et al. 2006; Hochstein and Ahissar 2002; Ullman 1995).

#### REFERENCES

- Bar M.** A cortical mechanism for triggering top-down facilitation in visual object recognition. *J Cogn Neurosci* 15: 600–609, 2003.
- Bar M, Kassam KS, Ghuman AS, Boshyan J, Schmidt AM, Dale AM, Hamalainen MS, Marinkovic K, Schacter DL, Rosen BR, Halgren E.** Top-down facilitation of visual recognition. *Proc Natl Acad Sci USA* 103: 449–454, 2006.
- Boyaci H, Fang F, Murray SO, Kersen D.** Response to lightness variations in early human visual cortex. *Curr Biol* 17: 989–993, 2007.
- Corbetta M, Shulman GL.** Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci* 3: 201–215, 2003.
- Datta R, DeYoe EA.** I know where you are secretly attending! The topography of human visual attention revealed with fMRI. *Vision Res* 49: 1037–1044, 2009.
- Deshpande G, Hu X, Stilla R, Sathian K.** Effective connectivity during haptic perception: a study using Granger causality analysis of functional magnetic resonance imaging data. *NeuroImage* 40: 1807–1814, 2008.
- Dolan RJ, Fink GR, Rolls E, Booth M, Holmes A, Frackowiak RSJ, Friston KJ.** How the brain learns to see objects and faces in an impoverished context. *Nature* 389: 596–599, 1997.

- Dougherty RF, Koch VM, Brewer AA, Fischer B, Modersitzki J, Wandell BA.** Visual field representations and locations of visual areas V1/2/3 in human visual cortex. *J Vis* 3: 586–598, 2003.
- Fang F, Boyaci H, Kersten D, Murray SO.** Attention-dependent representation of a size illusion in human V1. *Curr Biol* 18: 1707–1712, 2008b.
- Fang F, Kersten D, Murray SO.** Perceptual grouping and inverse fMRI activity patterns in human visual cortex. *J Vis* 8: 2.1–2.9, 2008a.
- Fischer J, Whitney D.** Attention narrows position tuning of population responses in V1. *Curr Biol* 19: 1356–1361, 2009.
- George N, Dolan RJ, Fink GR, Baylis GC, Russell C, Driver J.** Contrast polarity and face recognition in the human fusiform gyrus. *Nat Neurosci* 2: 574–580, 1999.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P.** Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293: 2425–2430, 2001.
- Heinen K, Jolij J, Lamme VA.** Figure–ground segregation requires two distinct periods of activity in V1: a transcranial magnetic stimulation study. *Neuroreport* 16: 1483–1487, 2005.
- Hochstein S, Ahissar M.** View from the top: hierarchies and reverse hierarchies in the visual system. *Neuron* 36: 791–804, 2002.
- Huang X, Paradiso MA.** V1 Response timing and surface filling-in. *J Neurophysiol* 100: 539–547, 2008.
- Hupé JM, James AC, Payne BR, Lomber SG, Girard P, Bullier J.** Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature* 394: 784–787, 1998.
- James TW, Humphrey GK, Gati JS, Servos P, Menon RS, Goodale MA.** Haptic study of three-dimensional objects activates extrastriate visual areas. *Neuropsychologia* 40: 1706–1714, 2002.
- Kamitani Y, Tong F.** Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8: 679–685, 2005.
- Kersten D, Mamassian P, Yuille A.** Object perception as Bayesian inference. *Annu Rev Psychol* 55: 271–304, 2004.
- Klein I, Dubois J, Mangin JF, Kherif F, Flandin G, Poline JB, Denis M, Kosslyn SM, Le Bihan D.** Retinotopic organization of visual mental images as revealed by functional magnetic resonance imaging. *Brain Res Cogn Brain Res* 22: 26–31, 2004.
- Knill DC, Pouget A.** The Bayesian brain: the role of uncertainty in neural coding and computation for perception and action. *Trends Neurosci* 27: 712–719, 2004.
- Knill DC, Richards W.** Editors. *Perception as Bayesian Inference*. Cambridge, UK: Cambridge Univ. Press, 1996.
- Kourtzi Z, Kanwisher N.** Representation of perceived object shape by the human lateral occipital complex. *Science* 293: 1506–1509, 2001.
- Lee TS.** A Bayesian framework for understanding texture segmentation in the primary visual cortex. *Vision Res* 5: 2643–2657, 1995.
- Lee TS, Mumford D.** Hierarchical Bayesian inference in the visual cortex. *J Opt Soc Am A Opt Image Sci Vis* 20: 1434–1448, 2003.
- Malach R, Reppas JB, Benson RR, Kwong KK, Jiang H, Kennedy WA, Ledden PJ, Brady TJ, Rosen BR, Tootell RB.** Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proc Natl Acad Sci USA* 92: 8135–8139, 1995.
- McKeeff TJ, Tong F.** The timing of perceptual decisions for ambiguous face stimuli in the human ventral visual cortex. *Cereb Cortex* 17: 669–678, 2006.
- Mohr HM, Linder NS, Linden DE, Kaiser J, Sireteanu R.** Orientation-specific adaptation to mentally generated lines in human visual cortex. *NeuroImage* 47: 384–391, 2009.
- Mumford D.** On the computational architecture of the neocortex II. *Biol Cybern* 66: 241–251, 1992.
- Murray SO, Boyaci H, Kersten D.** The representation of perceived angular size in human primary visual cortex. *Nat Neurosci* 9: 429–434, 2006.
- Murray SO, Kersten D, Olshausen BA, Schrater P, Woods DL.** Shape perception reduces activity in human primary visual cortex. *Proc Natl Acad Sci USA* 99: 15164–15169, 2002.
- Rao RP.** Bayesian inference and attentional modulation in the visual cortex. *Neuroreport* 16: 1843–1848, 2005.
- Ress D, Backus BT, Heeger DJ.** Activity in primary visual cortex predicts performance in a visual detection task. *Nat Neurosci* 3: 940–945, 2000.
- Roelfsema PR, Lamme VA, Spekreijse H.** Object-based attention in the primary visual cortex of the macaque monkey. *Nature* 395: 376–381, 1998.
- Serences JT.** Value-based modulations in human visual cortex. *Neuron* 60: 1169–1181, 2008.
- Serences JT, Boynton GM.** Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron* 55: 301–312, 2007a.
- Serences JT, Boynton GM.** The representation of behavioral choice for motion in human visual cortex. *J Neurosci* 27: 12893–12899, 2007b.
- Shuler MG, Bear MF.** Reward timing in the primary visual cortex. *Science* 311: 1606–1609, 2006.
- Stanley DA, Rubin N.** fMRI activation in response to illusory contours and salient regions in the human lateral occipital complex. *Neuron* 37: 323–331, 2003.
- Ullman S.** Sequence seeking and counter streams: a computational model for bidirectional information flow in the visual cortex. *Cereb Cortex* 1: 1–11, 1995.
- Williams MA, Baker CI, Op de Beeck HP, Shim WM, Dang S, Triantafylou C, Kanwisher N.** Feedback of visual object information to foveal retinotopic cortex. *Nat Neurosci* 11: 1439–1445, 2008.