

Neural representation of object orientation: A dissociation between MVPA and Repetition Suppression

Miles Hatfield, Michael McCloskey, Soojin Park *

Department of Cognitive Science, Johns Hopkins University, Baltimore MD 21218, USA

ARTICLE INFO

Article history:

Received 29 January 2016

Revised 13 April 2016

Accepted 21 May 2016

Available online 25 May 2016

Keywords:

Multi-voxel pattern analysis

Repetition Suppression

Lateral occipital complex

Object orientation

Representational similarity

ABSTRACT

How is object orientation represented in the brain? Behavioral error patterns reveal systematic tendencies to confuse certain orientations with one another. Using fMRI, we asked whether more confusable orientations are represented more similarly in object selective cortex (LOC). We compared two widely-used measures of neural similarity: multi-voxel pattern similarity (MVP-similarity) and Repetition Suppression. In LO, we found that multi-voxel pattern similarity was predicted by the confusability of two orientations. By contrast, Repetition Suppression effects in LO were unrelated to the confusability of orientations. To account for these differences between MVP-similarity and Repetition Suppression, we propose that MVP-similarity reflects the topographical distribution of neural populations, whereas Repetition Suppression depends on repeated activation of particular groups of neurons. This hypothesis leads to a unified interpretation of our results and may explain other dissociations between MVPA and Repetition Suppression observed in the literature.

© 2016 Published by Elsevier Inc.

Introduction

Representing an object's orientation is important in a variety of circumstances. For example, picking up a hammer requires an accurate representation of the hammer's orientation. In addition, the orientations of objects may be crucial for interpreting a scene (e.g., a chair upright versus on its side), and orientation can also influence judgments about an object's stability (Cholewiak et al., 2013) and center of mass (Barnett-Cowan et al., 2011).

Behavioral evidence concerning the tendency to confuse orientations with one another may offer insights into how the brain represents object orientation. In particular, mirror image views of objects are especially prone to confusion (e.g., Bradshaw et al., 1976; Corballis and Beale, 1976; Farrell, 1979; Mello, 1965; Sekuler and Houlihan, 1968; Sutherland, 1957; Wolff, 1971). Mirror-image confusion is typically conceptualized as a tendency to confuse an image with its *left–right reflection*—that is, with its reflection across a vertical axis (Fig. 1A). However, recent behavioral research (Gregory and McCloskey, 2010; Gregory et al., 2011) has challenged this conception, arguing that in most studies object-based axes have been confounded with extrinsic (i.e., non-object-based) vertical axes. In Fig. 1A, for example, the object's primary axis of elongation is aligned with the Extrinsic Vertical Axis. Given this confounding, confusions involving reflection across an

Extrinsic Vertical Axis (EVA reflection) could equally be described as reflections across the object's primary axis (OPA reflections; Fig. 1B).

Recent studies have dissociated EVA reflections from OPA reflections by presenting object stimuli at oblique orientations (Fig. 1C & D). Under these conditions, adult participants rarely made left–right EVA reflection errors (Fig. 1C), and instead made OPA reflections (Fig. 1D) more often than any other type of error (Gregory and McCloskey, 2010; Gregory et al., 2011). These results indicate the importance of distinguishing different mirror image relationships: OPA reflections are highly confusable, whereas EVA reflections are not. In addition, the findings raise new questions about the neural representation of object orientation. Does the greater behavioral confusability for OPA reflections relative to EVA reflections stem from the way object orientation is represented in the brain?

In this article we examine the neural representation of object orientation in the object-selective lateral occipital complex (LOC, Malach et al., 1995). In particular, we investigate the hypothesis that the behavioral confusability of orientations reflects the similarity of their representations in LOC. Recent work on the neural representation of mirror images is broadly consistent with this view. Neuroimaging studies have shown that cortical regions comprising the LOC represent mirror image views of objects and faces similarly (Dilks et al., 2011; Kietzmann et al., 2012; Axelrod and Yovel, 2012). These results are bolstered by comparable findings in Macaque IT, the putative homolog of LOC (Freiwald and Tsao, 2010; Rollenhagen and Olson, 2000). While potentially consistent with the view that confusable orientations are represented similarly in object-selective cortex, these studies either presented stimuli exclusively in “upright” orientations, or only used extrinsic (vertical or horizontal) axes of reflection. Such designs do

Abbreviations: MVP-similarity, Multi-voxel pattern similarity; OPA, Object Primary Axis reflection; EVA, Extrinsic Vertical Axis reflection.

* Corresponding author.

E-mail address: park@cogsci.jhu.edu (S. Park).




Fig. 1. Mirror images and axes of reflection. Object stimuli shown with Extrinsic Vertical Axis (EVA, dot endpoints) and Object Primary Axis (OPA, arrow endpoints) imposed. Top row: When stimuli are presented with EVA and OPA aligned, EVA reflections (A) are identical to OPA reflections (B). Bottom row: When presented at a tilt, EVA reflections (C) and OPA reflections (D) can be dissociated.

not allow differentiation of OPA and EVA reflections, which differ in their degree of behavioral confusability.

In two fMRI experiments we investigated the neural representation of OPA reflections, EVA reflections, and a range of other orientation relationships in object-selective cortex (LOC). Experiment 1 used Repetition Suppression (RS) to ask whether OPA reflections are represented more similarly than EVA reflections. Experiment 2 investigated the representation of object orientation across a wider range of orientation relationships, including but not limited to OPA and EVA reflections. Using a continuous carry-over design (Aguirre, 2007), we measured the similarity of neural representations using both RS and multi-voxel pattern similarity (MVP-similarity). We compared these neural similarity measures to the behavioral confusability rates for those same orientations.

The use of both RS and MVP-similarity measures also afforded an opportunity to ask whether these methods capture different aspects of neural similarity. Although both methods are frequently used to address questions of representational similarity, previous studies have found that RS and MVP do not always lead to the same conclusions (Drucker and Aguirre, 2009; Epstein and Morgan, 2012; Moore et al., 2013; Ward et al., 2013), raising pressing questions about what conclusions should be drawn in cases where the two methods diverge. In the *General discussion* we propose that RS and MVP-similarity are sensitive to different aspects of neural response, arguing that this hypothesis can explain our results as well as other dissociations between MVP and RS.

Experiment 1

This experiment used RS to ask whether object-selective cortex differentiates OPA and EVA reflections. On each trial, participants viewed two object images involving 1) an identical repetition; 2) an OPA reflection; 3) an EVA reflection; or 4) two different objects (Fig. 2A). All objects were presented in oblique orientations (Fig. 2B) so that OPA reflections were always distinct from EVA reflections. If the behavioral




Fig. 2. (A) Schematic illustration of the experimental conditions. (B) Each stimulus was presented in eight different oblique orientations across runs.

confusability of orientations is related to the degree of RS, we expect greater RS for the OPA condition than for the EVA condition.

Methods

Participants

Participants were 25 adults (15 males) between 18 and 33 years of age, with normal or corrected-to-normal vision. Three participants were removed due to our inability to localize any of the object-selective regions at $FDR < .05$, and 3 more were removed due to excessive head motion, leaving 19 participants for analysis.

Stimuli

Stimuli were color photographs of 60 isolated objects. All objects were poly-oriented (possessing no obvious canonical orientation), had a clear primary axis of elongation, and were bilaterally asymmetric. Most were tools or other manipulable objects. Objects were presented in 8 different oblique orientations (Fig. 2B) and subtended approximately 4.6° of visual angle. All stimuli were presented using the Psychophysics Toolbox 3 (Brainard, 1997).

Object orientation runs

Participants completed 10 object orientation runs, each consisting of 52 2-second trials, with 52 variable length (jittered 2–8 s) fixation trials in between (152 TRs total). On each trial two object images were presented for 500 ms each, with a 500 ms blank after each image (Fig. 3). The first image in a trial was followed by either the same image (Identical condition), a reflection across the object's primary axis (OPA condition), a reflection across an Extrinsic Vertical Axis (EVA condition), or a different object with its primary axis of elongation oriented the opposite way (Different condition); see Fig. 2A. Trials were counterbalanced so that within each run, each of the 60 different objects in the stimulus




Fig. 3. Trial structure for Experiment 1.

set appeared exactly once and, across runs, every object appeared in each condition the same number of times (each time in a different orientation). To maintain participants' attention, three additional (non-analyzed) trials with images of toy vehicles were included in each run, and participants pressed a button when they detected a vehicle.

Localizers and regions of interest

In an object localizer run, participants passively viewed 8 randomly ordered blocks of object images (different from those used in the main experiment) and 8 blocks of scrambled images (Park et al., 2007; Grill-Spector et al., 1999). Scrambled images were created by dividing the object images into a 16×16 square grid and randomly shuffling the squares based on eccentricity (Kourtzi and Kanwisher, 2000). In each block, 20 images were displayed for 600 ms each, followed by a 200 ms blank, with 10 s of rest (fixation) between blocks. Across the 7.1 minute run 160 objects and 160 scrambled images were presented.

Object-selective ROIs were functionally defined as regions in occipital and ventral temporal cortex that responded more strongly to objects than to scrambled objects ($FDR < .05$), with cluster threshold of 4 voxels (see Fig. 4). Object-selective LOC is divided into two subregions, the more posterior portion LO (lateral occipital) and more anterior pFs (posterior fusiform) (Grill-Spector et al., 1999). When the region of active voxels for the objects > scrambled images contrast was continuous and encompassed both LO and pFs, these ROIs were separated by increasing the threshold to $p < .00001$ ($N = 3$).

V1 was localized with a retinotopic mapping run. Participants fixated the center of the screen while a wedge-shaped checkerboard flashed. Five blocks of each condition (horizontal and vertical) were presented in an alternating pattern, each block lasting 12 s, followed by 12 s of fixation. The entire run lasted 4.2 min (126 TRs). The borders of early retinotopic regions were defined by a contrast of horizontal > vertical wedges at $FDR < .05$, after which V1 was manually drawn in to separate it from V2–V3 (Spiridon and Kanwisher, 2002; Park et al., 2007).

fMRI acquisition

Scanning was performed on a 3T Phillips scanner at the F.M. Kirby Center for functional brain imaging at the Kennedy Krieger Institute (Baltimore, MD). Functional images were acquired using a 32-channel head coil and a gradient echo single-shot echo planar imaging sequence [36 slices, repetition time (TR) = 2 s, echo time (TE) = 30 ms, flip angle = 70° , voxel size $2.5 \times 2.5 \times 2.5$ mm, 0.5 mm interslice gap]. Slices

were oriented along the transverse plane, covering from the most superior point of the brain to the bottom of the middle temporal lobe. High-resolution ($1 \times 1 \times 1$ mm voxel size) magnetization-prepared rapid acquisition images with gradient echo (MPRAGE) anatomical images were also acquired for each participant in order to overlay functional activity on the cortical surface.

Preprocessing and data analysis

Preprocessing with Brain Voyager QX 2.6.1 (Brain Innovation) included slice scan-time correction, linear trend removal, high-pass filtering and three-dimensional motion correction. Voxel timecourses were then z-transformed within each run. Motion parameters were included as regressors of no interest in all subsequent GLM analyses. No additional spatial or temporal smoothing was performed. All analyses were performed in native space (after ACPC alignment of functional and anatomical images).

After ROIs were defined for each participant, timecourses for each ROI for each participant were obtained by extracting the beta weights from a univariate GLM modeling the 10 contiguous timepoints (FIR model) for each condition starting with the stimulus onset. For each ROI the peak beta weights for each condition were identified by testing the timepoint with the highest average (numerical peak) across participants against all neighboring time points via two-tailed paired t-tests. For all ROIs, timepoint 5 (8 s after stimulus onset) was significantly higher than all other timepoints (all p 's < .01). All subsequent analyses are based on the beta values calculated at this timepoint.

Results

For each ROI (V1, LO, and pFs) potential hemispheric differences were assessed with a 2 (Hemisphere: Right vs. left) $\times$ 4 (Condition: Identical, OPA, EVA, Different) repeated-measures ANOVA. No main effects or interactions involving hemisphere were observed, so all subsequent analyses collapsed across hemispheres.¹ The average timecourses for LO, pFs, and V1 are shown in Fig. 5, with peak responses for LO and pFs shown in Fig. 6.

For all ROIs, repeated-measures ANOVAs carried out on the peak-response data for each condition (Identical, OPA, EVA, Different) revealed a significant effect of condition: LO, $F(3,54) = 16.4$, $p < .001$; pFs, $F(3,45) = 26.28$, $p < .001$, V1, $F(3,51) = 3.26$, $p = .0289$. Subsequent analyses explored the condition effects with paired t-tests (Bonferroni corrected for multiple comparisons).

Both LO and pFs showed RS: peak response was reduced in the Identical condition relative to the Different condition (LO: $t(18) = 6.26$, $p < .001$; pFs, $t(15) = 8.06$, $p < .001$). For V1 the Identical condition showed slightly lower peak response than the Different condition, but this difference was not significant ($t(17) = 1.20$, $p = .24$). This result is consistent with other studies failing to find RS to Identical repetitions in V1 (Ewbank et al., 2005; Grill-Spector et al., 1999; Sayres and Grill-Spector, 2008; Weiner et al., 2010). Subsequent analyses focused on the two object-selective regions, LO and pFs.

We first asked whether repetitions of the same object identity (irrespective of orientation) elicited RS. Specifically, we compared the conditions involving the same object in a different orientation (OPA and EVA conditions) to the Different object condition. In LO the response to OPA and EVA reflections did not differ from the response to a completely different object ($t(18) = 2.10$, $n.s.$; $t(18) = 1.36$, $n.s.$, respectively). In addition, both OPA and EVA reflections differed from the Identical condition, Identical vs. OPA: $t(18) = 5.23$, $p < .001$; Identical vs. EVA: $t(18) = 5.19$, $p < .001$, indicating a significant release from RS for orientation changes. Finally, OPA and EVA reflections were not different from one another $t(18) = .78$, $n.s.$




Fig. 4. Object-selective regions of interest in a representative participant. LO is shown in red on top, and pFs is shown in green below. Average voxel coordinates across participants (Talairach x, y, z): LO-L: -44.25 , -69.31 , -6.37 ; LO-R: 42.12 , -70.10 , -7.57 ; pFs-L: -35.58 , -48.80 , -16.41 ; pFs-R: 32.10 , -49.82 , -16.74 .

¹ In cases where LO and pFs were localized in only one hemisphere, we report the values obtained in that hemisphere.




Fig. 5. Hemodynamic time courses (beta values) for LO, pFs, and V1. Error bars are 1 SEM.

In contrast to LO, pFs showed sensitivity to object identity independent of orientation: repetition of the same object identity at a different orientation (OPA and EVA conditions) reduced the peak response relative to the Different object condition: Different vs. OPA: $t(15) = 5.49$, $p < .001$; Different vs. EVA: $t(15) = 4.04$, $p = .001$. An ROI (LO, pFs) $\times$ Condition (Ident, OPA, EVA, Diff) two-way repeated-measures ANOVA revealed no main effect of ROI, $F(1,15) = 1.46$, $p = 0.25$, but a significant interaction between condition and ROI, $F(3,45) = 10.53$, $p < .001$. The bulk of this effect (82% of the interaction variance) was captured by a Same- vs. Different-Object contrast [(Ident + OPA + EVA) – 3*Diff] $\times$ ROI (LO, pFs). This result indicates that LO and pFs differ specifically on whether repetition of an object identity (irrespective of orientation) induced RS, with pFs showing a strong RS effect with repetition of object identity (even with a changed orientation), and LO failing to show this effect.

However, pFs was not completely insensitive to orientation changes: EVA reflections induced a significantly higher response than Identical repetitions: Identical vs. EVA: $t(15) = 4.32$, $p < .001$. OPA reflections also induced a numerically greater response than Identical reflections in pFs. This difference approached but did not reach significance at the Bonferroni-corrected $p = .0083$ alpha-level: Identical vs. OPA: $t(15) = 2.66$, $p = .017$.

Finally, pFs, like LO, did not distinguish OPA and EVA reflections (OPA vs. EVA: $t(15) = 2.12$, *n.s.*). Although the mean response across participants was slightly lower for OPA than EVA reflections in pFs (Fig. 6), only 10 of the 19 individual participants showed this effect, with the remaining 9 showing a difference in the opposite direction.

Experiment 1 discussion

We analyzed RS effects in object-selective regions LO and pFs induced by sequential presentations of Identical orientations, OPA reflections, EVA reflections, or different objects. We predicted that the behaviorally more confusable OPA reflections would elicit a greater degree of RS than EVA reflections. However, this prediction was not confirmed in either LO or pFs. Neither region showed any difference in RS between OPA and EVA reflections.

In pFs, both OPA and EVA reflections induced RS relative to the Different object condition. These results are consistent with previous work showing RS to “mirror image reversals” in pFs (Dilks et al., 2011), as both OPA and EVA reflections are mirror images. Our results do not, however, directly support an interpretation on which RS for mirror images is related to their behavioral confusability: the current experiment distinguished highly confusable OPA reflections from rarely confused EVA reflections, and found that they did not differ in their degree of RS. Given the present results, RS to both OPA and EVA reflections in pFs may reflect a generally more orientation- and view-invariant representation of objects in this region (Eger et al., 2004; Grill-Spector et al., 1999; Vuilleumier et al., 2005), without this RS having anything in particular to do with mirror-image relationships or their behavioral confusability. Indeed, our RS results in pFs could largely be explained by sensitivity to repetition of object identity, present for Identical, OPA, and EVA trials but not for Different object trials.

LO, on the other hand, showed RS only for exact stimulus repetitions, a pattern of results that could potentially be due to retinotopic sensitivity in LO. However, previous studies have demonstrated RS to the same object despite shifts in retinal position or rotations in depth (Andersen et al., 2009; MacEvoy and Epstein, 2007; James et al., 2002; though see Kravitz et al., 2010), arguing against a strictly retinotopic representation. In addition, OPA reflections produce highly similar retinal images, yet did not induce any RS. Nonetheless, we attempt to control for the




Fig. 6. Peak responses in object-selective ROIs. Error bars are 1 SEM.

influence of retinotopic similarity while further exploring the neural representation of object orientation in Experiment 2.

Experiment 2

Experiment 1 focused on determining whether and to what extent OPA and EVA reflections elicit RS. The design was not, however, suitable for MVP-similarity analyses, the other method of interest in the present study. In Experiment 1 we could not obtain multi-voxel patterns for individual stimulus orientations, because two stimulus images were presented on each trial, and the estimated beta values therefore reflected the response to two stimuli. Experiment 2 was designed to allow MVP-similarity and RS analyses. We used a continuous carry-over design (Aguirre, 2007), presenting a counterbalanced sequence of 16 different orientations of the same object.² This design allowed us to measure the multi-voxel patterns produced by each of the 16 individual orientations, while simultaneously measuring the RS effect of each orientation on its successor.

We addressed two questions about the representation of object orientation, as revealed by both RS and MVP-similarity. First, we asked whether the neural representations for OPA reflections are more similar than representations for EVA reflections. We sought to replicate the results of Experiment 1, which found no differences in RS between OPA and EVA reflections, while also assessing whether MVP-similarity measures would cohere with the RS results. Second, we asked whether object-selective cortex is sensitive to the behavioral confusability of orientations in general, not just for OPA and EVA reflections. To address this question, we measured the extent to which neural similarity correlates with behavioral confusability across a range of different orientation relationships (including but not limited to OPA and EVA reflections). Finally, based on the results from both RS and MVP-similarity analyses, we discuss whether RS and MVP-similarity measure different aspects of neural similarity, and how the observed differences between methods yield insights into the neural representation of object orientation.

Methods

Participants

Twelve adults (5 males) ages 19–28 with normal or corrected-to-normal vision participated in the study. One participant was removed due to our inability to localize any object-selective regions, leaving 11 participants for analysis.

Stimuli

Stimuli were 16 orientations of a single object (an ice cream scoop) selected from the stimulus objects used in the Gregory and McCloskey (2010) behavioral study (Fig. 7A). The scoop was chosen because the behavioral confusability results for this object were highly representative of those for the stimulus set as a whole.

Continuous carry-over runs

Stimuli were presented in counterbalanced T111 sequences as part of a continuous carry-over design (Aguirre, 2007). This design renders the “direct effect” of stimulus presentation for each stimulus (underlying multi-voxel pattern analyses) orthogonal to the “carry-over effect” of the response of a previous stimulus on the response to the current stimulus (underlying RS analyses), allowing both analyses to be performed on the same data set. All participants viewed the same three T111 sequences, with order counterbalanced across participants. Each sequence consisted of 306 trials: 17 presentations of each of the 16 stimuli, plus 34 (17 × 2) presentations of a blank trial, which was doubled on each occurrence. Each sequence was divided into sets of 3 back-to-back runs

using the scan overlap technique (Aguirre, 2007) for a total of 9 runs per participant.

On each stimulus trial (Fig. 7B), participants viewed the object stimulus in a specific orientation for 1500 ms, followed by a 500 ms blank screen, followed by a dot that appeared in one of four locations relative to the previously presented object: in front of or behind the “mouth” of the scoop, or in front of or behind the end of the handle. Two of the four locations were defined as target locations (in front of the mouth, and behind the handle), and participants were instructed to simultaneously press the buttons in both hands any time the dot appeared in either target location. This task ensured that participants attended to the entire object.

Model similarity matrices

Two model similarity matrices were generated for comparison to the MVP-similarity data: 1) the behavioral confusability between orientations, and 2) the degree of pixel overlap between orientations.

Behavioral confusability model. Behavioral data from Gregory & McCloskey (2010, Experiment 1a) were used to generate a 16×16 confusion matrix encoding the rate of confusion errors between each pair of the 16 different orientations of an object. From this raw confusion matrix, we generated a behavioral confusability estimate for each orientation relationship (e.g., OPA reflection, EVA reflection, 45° rotation). All matrix cells that were instances of the same orientation relationship were averaged, and all of the cells were assigned the average value (e.g., confusion rates were averaged across all cells that represented an OPA reflection error, and then each was re-assigned the average value). This procedure yielded a confusion matrix with the same dimensions as our MVP-similarity matrices, allowing us to directly compare the two.

Pixel overlap model. To estimate the effect of low-level image similarity, we computed a pixel overlap similarity matrix encoding the degree of image overlap between each of the 16 stimulus images. For each image pair pixel overlap was defined as overlapping pixels / (total pixels / 2).

Localizers and regions of interest

In addition to the continuous carry-over runs, participants performed the same object-selective region and V1 localizers described in Experiment 1. To improve identification of object-selective ROIs, a second object localizer using new stimuli and a one-back task (3.6 min, 109 TRs) was also presented. When object-selective ROIs could be identified with both object localizers, the final ROI was defined as the intersection of the regions defined by each separate localizer; otherwise, a single localizer was used. ROI definition procedures were otherwise the same as in Experiment 1.

fMRI acquisition & preprocessing

The main experimental runs used a 3 s TR (TE = 30 ms) and acquired 47 slices, otherwise all parameters were the same as Experiment 1. For both RS and MVPA analyses, all pre-processing steps were the same as Experiment 1.

RS analyses

Main RS analyses. For purposes of comparing the present results to those of Experiment 1, we modeled the average response to a stimulus as a function of its predecessor, for trials in which the predecessor was an identical image, an OPA reflection of the current stimulus, or an EVA reflection. A nuisance predictor for button responses was also included. We fit a separate model to each participant's data to produce beta estimates for each participant for each condition and each ROI.

² The behavioral confusability of two orientations is only defined within an object, necessitating the use of a single object stimulus.




Fig. 7. Stimulus set and trial structure for continuous carry-over runs. Left: Participants viewed a continuous sequence of 16 different orientations of a single object stimulus. Right: On each trial, participants viewed an oriented stimulus while performing a dot detection task to force attention to the entire stimulus object (see main text for description of task).

In a second analysis not limited to Identical, OPA, and EVA orientation relationships, we modeled neural response as a function of the relative behavioral confusability between consecutive stimuli (as defined by the behavioral confusability model) as well as low-level pixel overlap between them (as defined by the pixel overlap model). Both behavioral confusability and pixel overlap models were transformed into distance matrices and then used as predictors in a GLM to model a continuous RS effect (Aguirre, 2007; Drucker and Aguirre, 2009).

Exploratory whole-brain analysis. To assess whether any regions outside of our ROIs showed RS, we performed a whole-brain analysis assessing RS to OPA reflections or EVA reflections. All individual participant timecourses were transformed into Talairach space, and a group-level random effects (RFX) GLM was performed on the group data. We then performed whole-brain main effect (OPA, EVA) and contrast (EVA–OPA) analyses.

MVP-similarity analyses

To estimate multi-voxel patterns for each of the 16 stimuli, voxel timecourses were modeled with a GLM including categorical predictors for each of the 16 stimuli and a nuisance predictor modeling button presses, producing an estimate of the activation at each voxel, for each stimulus, in each of our ROIs. Voxel responses were then mean-centered by subtracting the mean response across stimuli from each voxel's response for each stimulus (Drucker and Aguirre, 2009; Haxby, 2001). Within each ROI, the mean-centered voxel responses for each stimulus were then correlated with one another in a pairwise fashion to generate a 16×16 correlation matrix (henceforth “MVP-similarity matrix”).

Statistical significance was assessed via non-parametric permutation analyses in which the test statistic of interest was first calculated with the “correct” ordering of rows and columns of the MVP-similarity matrix, and was then resampled 10,000 times with replacement (Kriegeskorte et al., 2008) under permuted orderings of the rows and columns of the matrix to generate a sampling distribution of the test statistic. Statistical significance was established by comparing the correlation from the “correct” ordering to the rest of the sampling distribution. “True” correlations that fell in the top or bottom 1% of the permuted samples, as set by our significance criterion of $\alpha = .01$ (two-tailed), were deemed statistically significant.

Comparing OPA and EVA reflections: MVP-similarity. We first compared the correlation of multi-voxel patterns for Identical orientations, OPA reflections, and EVA reflections. To obtain an informative value for Identical orientations, we used the MVP correlation values from the split-half analysis procedure (see Supplementary Fig. S1). For each participant, we extracted the cell values from the MVP-similarity matrix

corresponding to correlations between Identical orientations, OPA reflections, and EVA reflections, and averaged them within a participant. To test these values for significance, we simulated the null hypothesis that the average correlation for Identical pairs, OPA reflections, and EVA reflections does not differ from what would be expected by a randomly-selected group of cells from the MVP-similarity matrix. See Supplementary Methods: Comparing OPA and EVA reflections: MVP-similarity for graphical representation and detailed description of procedure.

Relating MVP-similarity to behavioral confusability. We compared the MVP-similarity matrices from each participant to the two model similarity matrices (behavioral confusability and pixel overlap) via correlation and regression analyses. First, we computed the Pearson correlation between each participant's MVP-similarity matrix and the two model matrices and averaged them across participants, comparing these values to the distribution of correlations found based on 10,000 permutations with rows and columns shuffled. In the regression analysis, we entered each participant's MVP-similarity matrix into a regression model including behavioral confusability and pixel overlap similarity matrices as predictors (as well as a constant), generating beta weights for both models for each participant, which were then averaged across participants to produce a group measure. These group measures were compared to the distribution of values produced when rows and columns of both model similarity matrices were shuffled.

Results

RS analyses

In Experiment 1, highly confusable OPA reflections did not differ in degree of RS from rarely confused EVA reflections in either LO or pFs. To assess whether this finding would replicate in Experiment 2, we measured the response to a stimulus when it was an Identical repetition, an OPA reflection, or an EVA reflection of the previous stimulus (Fig. 8).

Repeated-measures ANOVAs revealed an effect of condition (Identical, OPA, EVA) in LO, $F(2,20) = 7.91$, $p = .003$, and in pFs, $F(2,20) = 4.12$, $p = .032$, but not in V1 $F(2,20) = 1.49$, $p = .25$. For LO, Identical repetitions induced a significantly lower BOLD response than either OPA or EVA reflections (OPA: $t(10) = 4.22$, $p = .002$; EVA: $t(10) = 3.104$, $p = .011$), but OPA and EVA conditions were not different from one another, $t(10) = .216$, $p = .83$. The same pattern held for pFs, although the reduction in response for Identical repeats did not survive multiple comparisons at the Bonferroni-corrected alpha level.

In an additional RS analysis, we assessed whether there was a continuous RS effect as a function of the behavioral confusability or pixel




Fig. 8. Average responses in each ROI to trials involving Identical repetitions, OPA reflections, and EVA reflections. Error bars are 1 SEM.

overlap of subsequent stimuli (Aguirre, 2007; Drucker and Aguirre, 2009). However, in all of the measured ROIs, neither the relative behavioral confusability nor degree of pixel overlap between subsequent stimuli modulated RS (all p 's > .05).

Finally, to ensure that our ROI-based approach did not exclude brain regions that may show an RS effect for OPA reflections, we conducted an exploratory whole-brain analysis. Even at a very liberal threshold of $p < .01$ (uncorrected), no regions showed significant response reduction to OPAs relative to the average response to any object stimulus, nor a significant difference between OPAs and EVAs.

Taken together, these results replicate the key finding of Experiment 1: for both LO and pFs, the degree of RS does not discriminate between OPA and EVA reflections. More generally, our analyses suggest that RS (in our ROIs as well as the rest of the brain) is insensitive to the behavioral confusability of orientations.

MVP-similarity analyses

After estimating voxel responses to each of the 16 stimuli and correlating them to generate 16×16 MVP-similarity matrices, we performed two initial analyses aimed at ensuring the quality of our MVP data (see Supplementary Information). These analyses confirmed that in both V1 and LO, stimuli produced reliable activation patterns across runs (Supplementary Analysis 1: Consistency of stimulus MVPs) and between-stimulus similarity relations were consistent across participants (Supplementary Analysis 2: Consistency of between-stimulus similarity relations). However, this was not true of pFs, which showed at-chance classification for stimuli and non-significant across-participant correlations. MVP-similarity analyses therefore targeted V1 and LO.

Our main analyses asked whether confusable orientations show similar multi-voxel patterns. First, we compared multi-voxel correlations for Identical orientations, OPA reflections, and EVA reflections. Second, we asked how well the behavioral confusability of orientations and the degree of pixel overlap between them can each explain the observed MVP correlations.

Comparing OPA and EVA reflections: MVP-similarity. To assess whether multi-voxel patterns were more similar for OPA reflections than EVA reflections in LO, we compared the average correlation between multi-voxel patterns for Identical orientations, OPA reflections, and EVA reflections. The mean correlation across participants for each orientation relationship is shown in Fig. 9.

Identical orientations produced the highest multi-voxel correlations ($R = 0.17$, 95% CI: [0.146; 0.191]), which were significantly higher than

OPAs (Identical vs. OPA, mean difference = 0.05, 95% CI: [0.012; 0.094]) and EVAs (mean difference = 0.20, 95% CI: [0.163; 0.245]). Most interestingly, OPA and EVA reflections were significantly different (mean difference = 0.15, 95% CI: [0.110; 0.191]): multi-voxel patterns were reliably more correlated for OPA reflections than for EVA reflections.

This result provides the first evidence of a neural representation consistent with the behavioral confusability of orientations: MVP correlations for OPA reflections are significantly higher than those for EVA reflections. However, this result alone does not establish that LO is sensitive to the behavioral confusability of orientations. OPA reflections have more similar retinal images than EVA reflections, suggesting that high MVP correlations for OPA reflections might be due to retinotopic sensitivity in LO (Silson et al., 2013).

To assess sensitivity to behavioral confusability while accounting for a possible effect of retinotopy, we compared the observed MVP-similarity measures to two quantitative model similarity matrices: one based on behavioral confusion errors between stimuli and one based on pixel overlap between stimuli. Simple pixel-based similarity measures provide good models of retinotopic regions such as V1 (Allred et al.,




Fig. 9. MVP-similarity (average MVP correlation) for Identical orientations, OPA reflections, and EVA reflections. Error bars represent standard error of the mean from permutation tests.




Fig. 10. Model and observed similarity matrices. Top row: model similarity matrices of interest. Bottom row: across-participant average MVP-similarity matrix (Pearson's R) for V1 and LO.

2005; Grill-Spector et al., 1999; Op de Beeck et al., 2001; Op de Beeck et al., 2008), and thus we used pixel-overlap to approximate retinotopic similarity.

Relating MVP-similarity to behavioral confusability. For both V1 and LO we computed the Pearson correlation between the MVP-similarity matrix and the pixel overlap and behavioral confusability model matrices (Fig. 10). MVP-similarity in V1 was significantly correlated with both model similarity matrices (behavioral confusability: $R = 0.45$, 95% CI = [0.261, 0.639]; pixel overlap: $R = 0.86$, 95% CI = [0.675, 1]). However, correlations with pixel overlap were significantly higher than with behavioral confusability (mean difference = -0.41 , 95% CI = [-0.600 ,

-0.227]). The high correlation (.86) between pixel overlap and MVP-similarity in V1 validates the pixel overlap measure, suggesting that it captures aspects of low-level stimulus similarity to which V1 is highly sensitive.

In LO, both behavioral confusability and pixel overlap models were significantly correlated with MVP-similarity (behavioral confusability: $R = 0.33$, 95% CI = [0.237, 0.432]; pixel overlap: $R = 0.10$, 95% CI = [0.002, 0.199]). However, the pattern was the reverse of that observed in V1: for LO, the correlation with the behavioral confusability model was significantly greater than the correlation with the pixel overlap model (mean difference = 0.23 , 95% CI = [0.137, 0.331].)




Fig. 11. Simultaneous regression of LO MVP-similarity matrix with behavioral confusability and pixel overlap matrices as predictors. Error bars represent 1 standard error of the mean from the permutation tests.

The modest correlation of $R = .10$ between MVP-similarity in LO and the pixel overlap model may reflect the moderate correlation between the behavioral confusability and pixel overlap models ($R = .50$). To assess the unique contribution of each model to predicting MVP-similarity in LO, we performed a simultaneous regression analysis predicting LO MVP correlations using behavioral confusability and pixel overlap models as predictors (Fig. 11).

In LO, behavioral confusability was a reliable predictor of MVP correlations ($\beta = 3.10$, 95% CI = [2.151, 4.045]), but pixel overlap was not ($\beta = -0.01$, 95% CI = [-0.202, 0.006]), a difference that was significant (β difference 95% CI = [2.193, 4.198]). These results suggest that the weak correlation between LO MVP correlations and pixel overlap was largely due to the shared variance between behavioral confusability and pixel overlap similarity matrices.

The shared variance between pixel overlap and behavioral confusability models may also account for the correlation between behavioral confusability and V1 ($R = .45$). To examine this possibility and expand our analysis to other visual regions, we performed the above regression analysis as part of an MVP-searchlight analysis (Rothlein and Rapp, 2014), searching across all object-responsive voxels in ventral and occipito-parietal cortex (see Supplementary Analysis 3: MVP-similarity searchlight analysis). The results showed that while MVP-similarity in V1 and surrounding retinotopic cortex were well-modeled by pixel overlap, behavioral confusability did not make a significant unique contribution to the regression model for these regions. In contrast, more anterior visual regions corresponding approximately to bilateral LO showed a significant unique effect of behavioral confusability, but no unique effect of pixel overlap, replicating our ROI-based regression result. Taken together with the previous analysis, these results confirm that in LO 1) OPA reflections elicit more similar multi-voxel patterns than EVA reflections and 2) this is attributable to a sensitivity to the behavioral confusability of orientations, and not to sensitivity to low-level image features.

These results also indicate that LO is sensitive to the relative behavioral confusability of orientations other than OPA and EVA reflections. In

Fig. 12 we plot the relative behavioral confusability for several different orientation relationships, based on Gregory and McCloskey (2010) data (Fig. 12A), and the average LO MVP correlations found for those same relationships in Experiment 2 (Fig. 12B). The pattern of high and low correlation is well matched by the pattern of high and low behavioral confusability not only for OPA and EVA reflections, but also across the whole set of orientation relationships.

MVP-similarity–RS relationship. We have shown that in LO, MVP-similarity is sensitive to the behavioral confusability of orientations, but RS is not. The observed difference between MVP-similarity and RS was most striking for OPA reflections: multi-voxel pattern correlations were significantly higher for OPA reflections than for EVA reflections, yet OPA reflections did not induce any detectable RS. In an additional RS analysis restricted to LO voxels that contributed most strongly to the positive multi-voxel pattern correlation for OPA reflections, we still failed to observe RS for OPAs (see Supplementary Analysis 4: RS computed over LO voxels contributing to MVP-similarity). The clear contrast between RS and MVP results suggests that these measures may reflect different aspects of the neural response to oriented object stimuli.

Experiment 2 discussion

We have shown that in LO, OPA reflections elicit more strongly correlated multi-voxel patterns than EVA reflections. This differential MVP-similarity was found to be specifically attributable to their differential behavioral confusability, and not to their differential retinotopic similarity as approximated by the pixel overlap similarity matrix.

The use of the pixel overlap similarity matrix also mitigates a potential concern about the generality of these results—specifically, the fact that we used a single object as a stimulus. Pixel overlap captures aspects of the visual similarity relations that would be highly specific to a particular object, yet this low-level visual




Fig. 12. MVP-similarity compared to the proportion of confusion errors for a range of orientation relationships. For both the model behavioral confusability matrix (top) and the observed LO MVP-similarity matrix (bottom), the cells corresponding to different orientation relationships were extracted and their values were averaged. OSA: object secondary axis reflection; EHA: extrinsic horizontal axis reflection.

similarity was not sufficient to explain LO's MVP-similarity results. Furthermore, our behavioral confusability similarity matrix was based on the data from a larger set of objects (Gregory and McCloskey, 2010), and the observed correlation with this model suggests that our stimulus object's similarity pattern is representative of the other objects in that study.

Perhaps the most striking result concerned the dissociation between RS and MVP-similarity. Whereas MVP-similarity in LO was significantly related to the confusability of orientations, we found no such relationship in RS analyses of the same region. Davis et al. (2014) recently showed that univariate measures (such as RS) are specifically sensitive to variance in effect size across participants, whereas MVPA measures are robust to such differences, potentially explaining some cases in which significant MVPA results are observed in the absence of RS. This would be a plausible explanation of the dissociation between MVP-similarity and RS observed here if all or a majority of our participants showed some trend toward greater RS for more confusable orientations, and simply varied in the size of this effect. However, this was not observed. Response to OPAs was virtually identical to response to EVAs in Experiment 2 across participants, and in Experiment 1 roughly half of participants showed a reduced response to OPA reflections relative to EVA reflections, whereas the other half showed the opposite ordering. This is the pattern of results expected if there were no RS effect at all, rather than a reliably present RS effect with variable effect size.

In LO, although no RS was observed for OPA and EVA reflections, we did observe RS for identical repetitions. However, in V1 even identical repetitions did not elicit RS, despite producing highly correlated multi-voxel patterns. Some previous studies have also failed to observe RS in V1 (Ewbank et al., 2005; Grill-Spector et al., 1999; Sayres and Grill-Spector, 2008; Weiner et al., 2010). A potential explanation is that RS for orientation in V1 is especially sensitive to timing parameters (Boynton and Finney, 2003; Fang et al., 2005). This raises the possibility that at a different ISI or with a prolonged adapting stimulus, we would have observed significant RS for identical repeats in V1, and that the duration of RS was simply not optimal for V1 in this experiment. While this is a plausible explanation, given the trend toward RS for identical repetitions in V1 (see Fig. 8), the failure to find a significant effect should not be over-interpreted.

General discussion

Previous studies finding similar neural representations for mirror images (Axelrod and Yovel, 2012; Dilks et al., 2011; Freiwald and Tsao, 2010; Kietzmann et al., 2012; Rollenhagen and Olson, 2000) raised the possibility that object-selective cortex is sensitive to the behavioral confusability of different orientations of an object. The present study explored this possibility systematically. Using RS (Experiments 1 and 2) and MVP-similarity (Experiment 2), we assessed the similarity of neural representations for two types of mirror image (OPA and EVA reflections) as well as other orientation relationships that vary in behavioral confusability.

MVP-similarity analyses in LO revealed a neural correlate of the actual mirror image confusion errors commonly made by participants (Gregory and McCloskey, 2010). Specifically, multi-voxel patterns were reliably more correlated for highly confusable OPA reflections than for rarely confused EVA reflections. This sensitivity to behavioral confusability extended to other orientation changes, and was robust even after accounting for similarity in low-level image properties.

RS analyses, however, showed a very different pattern. In Experiment 1, neural responses to OPA and EVA reflections in LO were not statistically distinguishable, and neither elicited any more RS than presentations of different objects. Experiment

2 again found no difference in RS between OPAs and EVAs in the same data for which MVP-similarity effects were observed. Even when a broader range of orientation changes was considered, no relationship was found between RS and behavioral confusability.

In the following discussion, we develop a working hypothesis to account for the seeming inconsistency between RS and MVP-similarity. To interpret our RS findings, we adopt the widely-accepted view that RS between two stimuli depends on an overlap in the neural populations responding to them (Grill-Spector and Malach, 2001). To interpret the MVP-similarity results, we posit that the correlation of multi-voxel patterns is sensitive to the anatomical clustering patterns of neural populations.

Relating RS and MVP-similarity: a working hypothesis

RS is sensitive to overlap of neural populations

Different models of RS posit different neural mechanisms to explain the reduced BOLD response observed for repeated stimuli (Grill-Spector et al., 2006; Gotts et al., 2012). However, all current models share the assumption that RS occurs only (although not necessarily always) when two stimuli activate some of the same neurons (Grill-Spector and Malach, 2001).³ This basic premise is sufficient for our interpretation of RS, which states that RS effects are sensitive to whether or not two stimuli activate overlapping sets of neurons.

MVP-similarity is sensitive to anatomical clustering of neural populations

Unlike RS, MVPA does not straightforwardly reflect whether two stimuli activate the same neural populations. Most multi-voxel analyses use the voxel's overall response to a stimulus (e.g. beta weight or t-value) to define the multi-voxel pattern for that stimulus. This voxel response is subject to the "averaging problem" often cited to motivate the RS approach (Grill-Spector and Malach, 2001; Malach, 2012). The averaging problem arises when a voxel responds similarly to two different stimuli, and the question is raised whether the same neurons are responding to both stimuli. Since the voxel's response reflects the activity of all of its constituent neurons, two different stimuli may elicit an equivalent response in a voxel by 1) activating the same neural population, or 2) activating distinct neural populations within that voxel. The averaging problem refers to the fact that, at the level of the voxel's response, these two situations cannot be distinguished.

Although the averaging problem is typically illustrated with a single voxel, it is of special relevance for multi-voxel analyses. The case in which many voxels respond similarly to two stimuli is exactly when a high multi-voxel correlation is observed. We illustrate this point schematically in Fig. 13, considering how the multi-voxel correlation between two stimuli varies based on whether they activate the same or different neural populations. We consider three different hypothetical cases of interest.

In hypothetical Case 1 (Fig. 13A), both Stimulus 1 and Stimulus 2 activate the same population of neurons, represented by the yellow dots

³ The key point on which models differ is precisely how the overlap in activated neural populations relates to a reduction in BOLD response. The dependency of RS on overlap in activated neurons is most clear in the Fatigue and Facilitation models (Grill-Spector et al., 2006). On these models, RS is due to reduction (Fatigue) or quickening (Facilitation) of response in re-activated neurons. On the Sharpening model, the first presentation of a stimulus is thought to activate both neurons that encode features critical to identification of a stimulus (identification-critical neurons), as well as many additional neurons that encode identification-irrelevant features (identification-irrelevant neurons). RS upon repetition results from repeated activation of the identification-critical neurons, which are thought to repeatedly inhibit the more numerous but less active identification-irrelevant neurons, lowering overall BOLD response. The Predictive Coding model proposed by Friston (2005) is somewhat more complex, but has been argued to constitute a higher-level description of processes that lead to RS through mechanisms described by Facilitation and Sharpening models (Friston, 2012; Henson, 2012; Desimone, 1996).




Fig. 13. MVP-similarity and hypothetical underlying neuronal responses. Yellow dots represent active neurons responding above baseline. Blue dots represent inactive neurons responding at baseline. (A) In Case 1, two stimuli activate the same population of neurons, producing a similar multi-voxel pattern leading to high multi-voxel correlations. (B) In Case 2, the neuronal populations responding to each stimulus are "clustered" or reliably nearby, activating the same voxels and thereby eliciting similar multi-voxel patterns. (C) In Case 3, the two stimuli activate distinct neurons that are not clustered, activating different voxels and leading to low multi-voxel correlations.

in Voxel 1. Considering the voxel responses that would be produced, both stimuli would activate Voxel 1 to some extent, but not Voxel 2. As a result, the pattern of responses across voxels produced by Stimulus 1 (high response in Voxel 1, and a low/baseline response in Voxel 2) is similar to that produced by Stimulus 2, and the multi-voxel correlation between them would be relatively high.

In Case 2 (Fig. 13B), Stimulus 1 and Stimulus 2 activate different sets of neurons, but those neurons are anatomically clustered so as to be sampled by the same voxel (Voxel 1). Again, both stimuli would activate Voxel 1 but not Voxel 2, eliciting a high multi-voxel correlation. This case demonstrates the averaging problem as it applies to multi-voxel analyses: even when two stimuli activate different neural populations, they may produce correlated multi-voxel patterns.

Critically, however, this effect requires that the different neural populations are situated in the same measured voxel. In Case 3 (Fig. 13C), Stimulus 1 and Stimulus 2 activate distinct neural populations that, in addition, are located in different voxels. The two stimuli would now activate different voxels, producing different patterns of responses and a low multi-voxel correlation.

This example is highly simplified, but nevertheless illustrates a key point. The observation of a high multi-voxel correlation for two stimuli is consistent with either Case 1 or Case 2: stimuli may activate the same (Case 1), or distinct-yet-reliably-clustered (Case 2) neural populations. A key advantage of supplementing MVP-similarity results with RS analyses is precisely the ability to disambiguate these two cases. Given a high MVP correlation, RS provides an independent method for determining whether the correlation is due to activation of overlapping versus distinct neural populations. Further, if it can be determined that high multi-voxel correlations arise from activation of distinct neural populations (i.e., that Case 2 holds), we gain insight into the topographical distribution of neural populations in cortex—information not necessarily provided by RS analyses alone. We demonstrate this point by applying our working hypothesis to the current results to elucidate how object orientation is represented in LO.

Object orientation in LO: the working hypothesis applied to the present results

Consider first our results for Identical orientations. The multi-voxel patterns for Identical orientations (i.e., the patterns for the same stimulus across runs) were correlated, a result that may reflect either activation of the same (Case 1) or distinct-yet-clustered (Case 2) neural populations. RS analyses revealed that successive presentations of Identical orientations elicited RS, supporting the Case 1 scenario (Fig. 13A).

The results for OPA reflections were importantly different. Like Identical orientations, OPA reflections elicited significant multi-voxel correlations, consistent with either Case 1 or Case 2. However, OPA reflections did not induce RS, suggesting that different neural populations are activated by each stimulus. Together, these results suggest that distinct-yet-reliably-clustered neural populations respond to OPA reflections, similar to Case 2 (Fig. 13B).

Finally, multi-voxel patterns for EVA reflections were not correlated, and also did not elicit RS. These results are consistent with Case 3 (Fig. 13C), suggesting that EVA reflections activate distinct neurons that also are not reliably clustered in the same voxels.

In sum, the difference between OPA and EVA reflections seems to be topographical: OPA reflections are represented by distinct-yet-clustered neural populations, whereas EVA reflections are not.⁴ In the final section, we discuss the functional relevance of neural clustering. First, we discuss how our working hypothesis relates to previous work on the MVPA-RS relationship, and discuss its critical assumptions.

Relations to previous work

The working hypothesis presented here builds on an account discussed by Drucker and Aguirre (2009) concerning the neural basis of MVPA. Drucker and Aguirre suggest that significant MVPA decoding for a stimulus indicates the presence of a "coarse spatial code" for that stimulus. However, the nature of this code is not discussed in detail so as to generalize beyond the context of object shape representation. In contrast, we have offered a specific and generalizable hypothesis about the neural basis of MVP-similarity, on which it is sensitive to the anatomical clustering patterns of neural populations. We have shown

⁴ Although not directly relevant to the present results, it is also of interest to understand how one might observe RS without MVP-similarity (Drucker and Aguirre, 2009, Ward et al., 2013). To explain how stimuli represented by distinct neural populations can produce correlated MVPs, we emphasize that those distinct neural populations must be *reliably* clustered. This emphasis on "reliably" is important: for two stimuli to produce correlated MVPs, they must elicit similar responses across many voxels. If the neural populations activated by two stimuli were clustered only in a small portion of cortex sampled by a few voxels, with no clustering relationship holding for the majority of measured voxels, the fact that they elicit similar responses in that small subset of voxels may not outweigh the fact that in most voxels, they elicit unrelated responses.

Similar considerations can explain how stimuli that activate the same neurons, and potentially elicit RS, can nonetheless fail to produce correlated MVPs. For example, if two stimuli activate the same neurons in a region encompassed by 10 voxels, but an MVP-similarity analysis includes those voxels as well as 90 others not involved in representing the stimulus, the similarity of response in those 10 voxels may be "washed out" by the inclusion of those additional 90 voxels in the analysis.

how different clustering patterns may lead to different MVPA results in a manner that is not tied to any particular stimulus type or property.

Other work has discussed additional differences between RS and MVPA consistent with, but not explicitly described by our working hypothesis. Epstein and Morgan (2012) discuss the account provided by Drucker and Aguirre (2009) as well as two additional hypotheses on which MVPA differs from RS 1) in its sensitivity to neural outputs as opposed to inputs, and 2) in virtue of reflecting a bottom-up instead of a top-down component of stimulus representation. As we understand them, these hypotheses are not mutually exclusive with the present account. Rather, they describe additional dimensions along which RS and MVPA could potentially differ. In a similar vein, Davis et al. (2014) showed that univariate measures like RS are more sensitive to variance in effect size across participants, whereas MVPA measures are robust to such differences. We have argued in the *Experiment 2 discussion* that this is unlikely to explain the present results, although such factors could potentially explain other dissociations between RS and MVPA.

Assumptions of the working hypothesis

In interpreting RS effects, we have adopted the widely-shared assumption that RS is sensitive to whether two stimuli activate overlapping neural populations. However, this assumption may be challenged. For instance, RS effects are known to be modulated by experimental factors such as the lag between stimuli (Epstein et al., 2008). This would seem to create difficulties for attempts to read properties of neural representation directly from RS effects. We agree that the effect of lag should not be ignored—as discussed in the *Experiment 2 discussion*, we believe that such factors may explain our failure to find significant RS for V1. However, the fact that RS is sensitive to such experimental factors is not inconsistent with the claim that neural overlap is a *necessary* condition for observing RS, the view we advocate here. Two stimuli may activate the same neurons but fail to show RS for other reasons, and one must keep this in mind when interpreting failures to observe RS in specific cases. In the present case, given that LO showed RS to identical repetitions at the present lag, our failure to find RS for OPA reflections and EVA reflections cannot readily be explained by a general failure to observe any RS at this lag.

Our working hypothesis also makes several simplifying assumptions concerning MVP-similarity. In our examples, each stimulus activated a similar number of neurons, but in practice this cannot generally be assumed. It will be important to understand how multi-voxel correlations behave as a function of differences in the number of neurons that two stimuli activate. The scale of the clustering patterns is also important. Specifically, with sufficiently small clustering patterns relative to voxel size, even comparatively distant neural populations may fall within the same voxel just as often as clustered populations. Under such circumstances, no differences would be detectable in the magnitude of multi-voxel correlations for distant as compared with reliably clustered neural populations, suggesting that MVP-similarity analyses may not be able to detect clustering patterns at all scales. Although further development is needed, we believe that this working hypothesis provides a simple, clear and plausible account of an important difference between MVPA and RS.

The functional relevance of neural clustering

We have suggested that analyses of multi-voxel similarity patterns, especially when used in conjunction with RS analyses, may provide a tool for revealing aspects of neural topography. Using this approach, we have argued that neurons responding to OPA reflections are more reliably clustered together than neurons responding to EVA reflections. These results join several recent findings showing that perceptually similar stimuli elicit correlated multi-voxel patterns (Drucker and Aguirre, 2009; Haushofer et al., 2008; Mur et al., 2013; Op de Beeck et al., 2008), raising the intriguing hypothesis that these similarity

patterns arise because perceptually similar stimuli are often represented by clustered or anatomically nearby neural populations.

Potentially illuminating hypotheses about the functional relevance of neural clustering patterns are found in the literature on cortical columns (Mountcastle, 1978; Mountcastle, 1997; Rockland and Ichinohe, 2004; Tanaka, 2003). According to the “common input” hypothesis, neurons may cluster together because they receive inputs from the same source, presumably from neurons in an earlier processing stage. Given that OPA reflections elicit similar retinal projections, it is plausible that the neurons that represent a given orientation and those that represent its OPA reflection would receive input from some of the same early visual neurons. EVA reflections are far more variable in this respect, often eliciting quite different retinal projections. Thus, common inputs from early retinotopic regions may partly explain LO's clustering patterns. However, this account only readily explains clustering patterns for stimuli that produce similar retinal projections, yet correlated multi-voxel patterns have been observed for stimuli that are similar in more abstract ways (e.g., in terms of being animate versus inanimate, Mur et al., 2013). In addition, as we have argued above, LO's multi-voxel similarity patterns cannot be entirely accounted for by retinotopy.

The “common output” hypothesis, in contrast, suggests that cells cluster based on where they send their output. OPA reflections plausibly elicit similar action affordances (for example, the positioning of the hand in order to grasp and interact with an object will often be similar for a given orientation and its OPA reflection), suggesting a common output signal to neurons involved in action planning (Culham and Valyear, 2006). EVA reflections, on the other hand, would often require quite different grasping postures. Recent work finds that reciprocal connections between ventral and dorsal streams are widespread in the Macaque (see Kravitz et al., 2013 for a review), suggesting that connections between LO and dorsal stream regions involved in action are not implausible.

Finally, it remains an open possibility that clustering patterns are epiphenomenal, serving no functional purpose (Horton and Adams, 2005), although on this view the relationship between the behavioral confusability of OPA reflections and their clustering patterns would be somewhat mysterious. Of course, these interpretations are purely speculative and require measures of neural clustering and connectivity at a much finer scale to be seriously tested. Nonetheless, we believe that questions about neural topography are of great interest to neuroscience, and that by combining two common fMRI methods, researchers can address them.

Conclusion

The present study found a clear dissociation between Repetition Suppression (RS) and multi-voxel pattern similarity (MVP-similarity) in region LO. This dissociation can be explained by a simple working hypothesis: MVP-similarity is sensitive to whether stimuli activate anatomically clustered neuronal populations, whereas RS is sensitive to whether they activate the same neuronal populations. This hypothesis not only accounts for our results but may also account for other apparent discrepancies between RS and MVPA reported in the literature.

Acknowledgements

This work was supported by research funding from the Johns Hopkins University to S.P. We thank David Rothlein for providing the code used for the MVPA searchlight analysis, Jung Uk Kang and Harry Ngai for their assistance with data collection and analysis, and Kristen Johannes for constructive comments on the paper.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <http://dx.doi.org/10.1016/j.neuroimage.2016.05.052>.

References

- Aguirre, G.K., 2007. Continuous carry-over designs for fMRI. *NeuroImage* 35 (4), 1480–1494.
- Allred, S., Liu, Y., Jagadeesh, B., 2005. Selectivity of inferior temporal neurons for realistic pictures predicted by algorithms for image database navigation. *J. Neurophysiol.* 94 (6), 4068–4081.
- Andersen, D.R., Vinberg, J., Grill-Spector, K., 2009. The representation of object viewpoint in human visual cortex. *NeuroImage* 45 (2), 522–536.
- Axelrod, V., Vovel, G., 2012. Hierarchical processing of face viewpoint in human visual cortex. *J. Neurosci.* 32 (7), 2442–2452.
- Barnett-Cowan, M., Fleming, R.W., Singh, M., Bühlhoff, H.H., 2011. Perceived object stability depends on multisensory estimates of gravity. *PLoS One* 6 (4), e19289.
- Boynton, G.M., Finney, E.M., 2003. Orientation-specific adaptation in human visual cortex. *J. Neurosci.* 23, 8781–8787.
- Bradshaw, J., Bradley, D., Patterson, K., 1976. The perception and identification of mirror-reversed patterns. *Q. J. Exp. Psychol.* 28 (2), 221–246.
- Brainard, D.H., 1997. The psychophysics toolbox. *Spat. Vis.* 10, 433–436.
- Cholewiak, S.A., Fleming, R.W., Singh, M., 2013. Visual perception of the physical stability of asymmetric three-dimensional objects. *J. Vis.* 13 (4), 12.
- Corballis, M.C., Beale, J.L., 1976. *The Psychology of Left and Right*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Culham, J.C., Valyear, K.F., 2006. Human parietal cortex in action. *Curr. Opin. Neurobiol.* 16 (2), 205–212.
- Davis, T., LaRocque, K.F., Mumford, J.A., Norman, K.A., Wagner, A.D., Poldrack, R.A., 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.
- Desimone, R., 1996. Neural mechanisms for visual memory and their role in attention. *Proc. Natl. Acad. Sci.* 93 (24), 13494–13499.
- Dilks, D.D., Julian, J.B., Kubilius, J., Spelke, E.S., Kanwisher, N., 2011. Mirror-image sensitivity and invariance in object and scene processing pathways. *J. Neurosci.* 31 (31), 11305–11312.
- Drucker, D.M., Aguirre, G.K., 2009. Different spatial scales of shape similarity representation in lateral and ventral LOC. *Cereb. Cortex* 19 (10), 2269–2280.
- Eger, E., Henson, R.N.A., Driver, J., Dolan, R.J., 2004. BOLD repetition decreases in object-responsive ventral visual areas depend on spatial attention. *J. Neurophysiol.* 92 (2), 1241–1247.
- Epstein, R.A., Morgan, L.K., 2012. Neural responses to visual scenes reveals inconsistencies between fMRI adaptation and multivoxel pattern analysis. *Neuropsychologia* 50 (4), 530–543.
- Epstein, R.A., Parker, W.E., Feiler, A.M., 2008. Two kinds of fMRI repetition suppression? Evidence for dissociable neural mechanisms. *J. Neurophysiol.* 99 (6), 2877–2886.
- Ewbank, M.P., Schluppeck, D., Andrews, T.J., 2005. fMR-adaptation reveals a distributed representation of inanimate objects and places in human visual cortex. *NeuroImage* 28 (1), 268–279.
- Fang, F., Murray, S.O., Kersten, D., He, S., 2005. Orientation-tuned fMRI adaptation in human visual cortex. *J. Neurophysiol.* 94, 4188–4195.
- Farrell, W.S., 1979. Coding left and right. *J. Exp. Psychol. Hum. Percept. Perform.* 5 (1), 42–51.
- Freiwald, W.A., Tsao, D.Y., 2010. Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science* 330 (6005), 845–851.
- Friston, K., 2005. A theory of cortical responses. *Philos. Trans. R. Soc. B Biol. Sci.* 360 (1456), 815–836.
- Friston, K., 2012. Predictive coding, precision and synchrony. *Cogn. Neurosci.* 3 (3–4), 238–239.
- Gotts, S.J., Chow, C.C., Martin, A., 2012. Repetition priming and repetition suppression: a case for enhanced efficiency through neural synchronization. *Cogn. Neurosci.* 3 (3–4), 227–237.
- Gregory, E., McCloskey, M., 2010. Mirror-image confusions: implications for representation and processing of object orientation. *Cognition* 116 (1), 110–129.
- Gregory, E., Landau, B., McCloskey, M., 2011. Representation of object orientation in children: evidence from mirror-image confusions. *Vis. Cogn.* 19 (8), 1035–1062.
- Grill-Spector, K., Malach, R., 2001. fMR-adaptation: a tool for studying the functional properties of human cortical neurons. *Acta Psychol.* 107 (1), 293–321.
- Grill-Spector, K., Kushnir, T., Edelman, S., Avidan, G., Itzhak, Y., Malach, R., 1999. Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron* 24 (1), 187–203.
- Grill-Spector, K., Henson, R., Martin, A., 2006. Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn. Sci.* 10 (1), 14–23.
- Haushofer, J., Livingstone, M.S., Kanwisher, N., 2008. Multivariate patterns in object-selective cortex dissociate perceptual and physical shape similarity. *PLoS Biol.* 6 (7), e187.
- Haxby, J.V., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293 (5539), 2425–2430.
- Henson, R.N., 2012. Repetition accelerates neural dynamics: in defense of facilitation models. *Cogn. Neurosci.* 3 (3–4), 240–241.
- Horton, J.C., Adams, D.L., 2005. The cortical column: a structure without a function. *Philos. Trans. R. Soc. B Biol. Sci.* 360 (1456), 837–862.
- James, T.W., Humphrey, G.K., Gati, J.S., Menon, R.S., Goodale, M.A., 2002. Differential effects of viewpoint on object-driven activation in dorsal and ventral streams. *Neuron* 35 (4), 793–801.
- Kietzmann, T.C., Swisher, J.D., Konig, P., Tong, F., 2012. Prevalence of selectivity for mirror-symmetric views of faces in the ventral and dorsal visual pathways. *J. Neurosci.* 32 (34), 11763–11772.
- Kourtzi, Z., Kanwisher, N., 2000. Cortical regions involved in perceiving object shape. *J. Neurosci.* 20 (9), 3310–3318.
- Kravitz, D.J., Kriegeskorte, N., Baker, C.I., 2010. High-level visual object representations are constrained by position. *Cereb. Cortex* 20 (12), 2916–2925.
- Kravitz, D.J., Saleem, J.S., Baker, C.I., Ungerleider, L.G., Mishkin, M., 2013. The ventral visual pathway: an expanded neural framework for the processing of object quality. *Trends Cogn. Sci.* 17 (1), 26–49. <http://dx.doi.org/10.1016/j.tics.2012.10.011>.
- Kriegeskorte, N., Mur, M., Bandettini, P.A., 2008. Representational similarity analysis – connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* 2 (4), 1–28.
- MacEvoy, S.P., Epstein, R.A., 2007. Position selectivity in scene- and object-responsive occipitotemporal regions. *J. Neurophysiol.* 98 (4), 2089–2098.
- Malach, R., 2012. Targeting the functional properties of cortical neurons using fMRI-adaptation. *NeuroImage* 62 (2), 1163–1169.
- Malach, R., Reppas, J.B., Benson, R.R., Kwong, K.K., Jiang, H., Kennedy, W.A., ... Tootell, R.B., 1995. Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proc. Natl. Acad. Sci.* 92 (18), 8135–8139.
- Mello, N.K., 1965. Interhemispheric reversal of mirror-image oblique lines after monocular training in pigeons. *Science* 148 (3667), 252–254.
- Moore, K.S., Yi, D.-J., Chun, M., 2013. The effect of attention on repetition suppression and multivoxel pattern similarity. *J. Cogn. Neurosci.* 25 (8), 1305–1314.
- Mountcastle, V.B., 1978. In: Edelman, G.M., Mountcastle, V.B. (Eds.), *An Organizing Principle for Cerebral Function: The Unit Module and the Distributed System*. The mindful brain, pp. 1–50.
- Mountcastle, V.B., 1997. The columnar organization of the neocortex. *Brain* 120 (4), 701–722.
- Mur, M., Meys, M., Bodurka, J., Goebel, R., Bandettini, P.A., Kriegeskorte, N., 2013. Human object-similarity judgments reflect and transcend the primate-IT object representation. *Front. Psychol.* 4 (128), 1–22.
- Op de Beeck, H.P., Torfs, K., Wagemans, J., Vogels, R., 2001. Inferotemporal neurons represent low-dimensional configurations of parameterized shapes. *Nat. Neurosci.* 4 (12), 1244–1252.
- Op de Beeck, H.P., Torfs, K., Wagemans, J., 2008. Perceived shape similarity among unfamiliar objects and the organization of the human object vision pathway. *J. Neurosci.* 28 (40), 10111–10123.
- Park, S., Intraud, H., Yi, D.-J., Widders, D., Chun, M.M., 2007. Beyond the edges of a view: boundary extension in human scene-selective visual cortex. *Neuron* 54 (2), 335–342.
- Rockland, K., Ichinohe, N., 2004. Some thoughts on cortical minicolumns. *Exp. Brain Res.* 158 (3).
- Rollenhagen, J.E., Olson, C.R., 2000. Mirror-image confusion in single neurons of the macaque inferotemporal cortex. *Science* 287 (5457), 1506–1508.
- Rothlein, D., Rapp, B., 2014. The similarity structure of distributed neural responses reveals the multiple representations of letters. *NeuroImage* 89, 331–344.
- Sayres, R., Grill-Spector, K., 2008. Relating retinotopic and object-selective responses in human lateral occipital cortex. *J. Neurophysiol.* 100 (1), 249–267.
- Sekuler, R.W., Houlihan, K., 1968. Discrimination of mirror-images: choice time analysis of human adult performance. *Q. J. Exp. Psychol.* 20 (2), 204–207.
- Silson, E.H., McKeefry, D.J., Rodgers, J., Gouws, A.D., Hymers, M., Morland, A.B., 2013. Specialized and independent processing of orientation and shape in visual field maps LO1 and LO2. *Nat. Neurosci.* 16 (3), 267–269.
- Spiridon, M., Kanwisher, N., 2002. How distributed is visual category information in human occipito-temporal cortex? An fMRI study. *Neuron* 35 (6), 1157–1165.
- Sutherland, N.S., 1957. Visual discrimination of orientation and shape by the octopus. *Nature* 179, 11–13.
- Tanaka, K., 2003. Columns for complex visual object features in the inferotemporal cortex: clustering of cells with similar but slightly different stimulus selectivities. *Cereb. Cortex* 13 (1), 90–99.
- Vuilleumier, P., Schwartz, S., Duhaux, S., Dolan, R.J., Driver, J., 2005. Selective attention modulates neural substrates of repetition priming and “implicit” visual memory: suppressions and enhancements revealed by fMRI. *J. Cogn. Neurosci.* 17 (8), 1245–1260.
- Ward, E.J., Chun, M.M., Kuhl, B.A., 2013. Repetition suppression and multi-voxel pattern similarity differentially track implicit and explicit visual memory. *J. Neurosci.* 33 (37), 14749–14757.
- Weiner, K.S., Sayres, R., Vinberg, J., Grill-Spector, K., 2010. fMRI-adaptation and category selectivity in human ventral temporal cortex: regional differences across time scales. *J. Neurophysiol.* 103 (6), 3349–3365.
- Wolff, P., 1971. Mirror-image confusability in adults. *J. Exp. Psychol.* 91 (2), 268.