SUPPLEMENTARY INFORMATION

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

Miles Hatfield, Michael McCloskey, & Soojin Park

SUPPLEMENTARY ANALYSES

Supplementary Analysis 1: Consistency of stimulus MVPs.

Figure S1: Data-splitting procedure

Figure S2: Average pairwise correlation-based classification

Supplementary Analysis 2: Consistency of between-stimulus similarity relations.

Figure S3: Consistency of between-stimulus similarity relations across participants.

Supplementary Analysis 3: MVP-similarity searchlight analysis.

Figure S4: MVP-similarity searchlight analysis results.

Supplementary Analysis 4: RS computed over LO voxels contributing to MVP-similarity

Figure S5: RS to OPAs is not related to similarity of voxel response to OPAs.

SUPPLEMENTARY METHODS

Comparing OPA and EVA reflections: MVP-similarity

Figure S6: Steps of Comparing OPA and EVA reflections: MVP-similarity analysis

Supplementary Analysis 1: Consistency of stimulus MVPs.

If multi-voxel patterns contain stimulus information, within-stimulus correlations should be higher than across-stimulus correlations. To assess this, we implemented a standard across-run correlation-based classification analysis (Haxby, 2001; Epstein & Morgan, 2012) modified to preserve the counterbalancing of the continuous carry-over design (Aguirre, 2007).

First, we generated MVP correlation matrices that represented the correlation of stimulus patterns across runs. For each participant we iteratively split the entire dataset in two and computed a separate GLM from each part of the data. Beta values from each part of the data were then mean-centered, correlated with those based on the remaining set of the data, and averaged together to produce the final correlation map for that ROI for that participant. See Figure S1 for a schematic representation of the procedure.

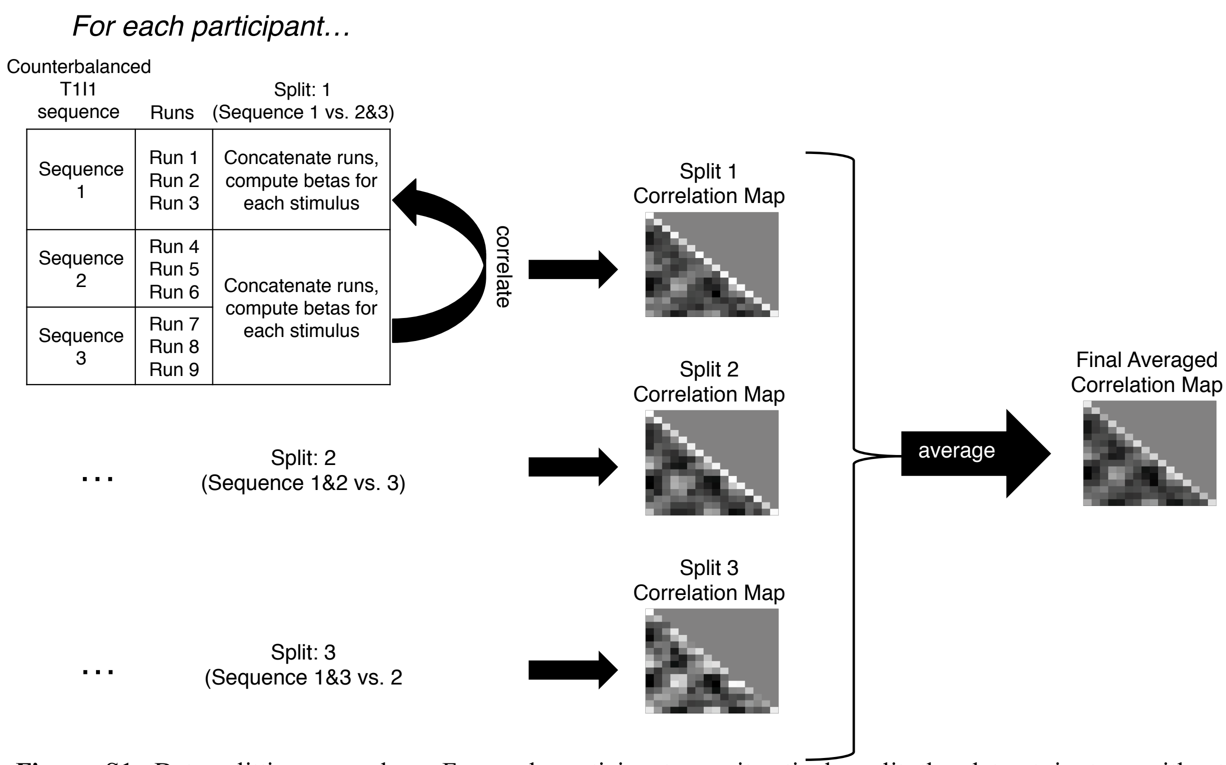

Figure S1. Data-splitting procedure. For each participant, we iteratively split the dataset in two without separating runs that together formed a fully counterbalanced T1I1 sequence: Split 1: Runs 1-3 vs. Runs 4-9; Split 2: Runs 1-6 vs. Runs 7-9; Split 3: Runs 1-3 & 7-9 vs. Runs 4-6. For each split, we computed a separate GLM from each part of the data (e.g. for Split 1, separate GLMs were computed from Runs 1-3, and Runs 4-9). Beta values from each member of a split were correlated with its corresponding member (e.g. for Split 1, betas estimated from Runs 1-3 vs. were correlated with betas from Runs 4-9, etc.) to produce a correlation map for that split. Averaging the 3 resulting correlation maps, we produce the final correlation map for that ROI for that participant.

For each of the 16 stimuli, the within-stimulus MVP correlation was compared to the between-stimulus correlations. A stimulus pattern was considered correctly classified (and assigned a value of 1) if the within-stimulus correlation for that stimulus was higher than the average across-stimulus correlation, otherwise it was assigned a zero (thus, given pairwise comparisons chance performance = 50%). Performance across all 16 stimuli was then averaged and tested against 50% chance with a one-tailed t-test (Epstein & Morgan, 2012).

Across-voxel patterns in V1 reliably discriminated between stimuli, with a classification performance of 96% ($t(9) = 48.67$, $p < .001$), Fig. S2. Both LO-L and R were also reliably above chance performance ($t(10) = 5.76$, $p < .001$; $t(9) = 4.31$, $p = .001$, respectively). However, neither pFs-L nor pFs-R showed above-chance classification, ($t(10) = 1.46$, $p=.13$; $t(7) = 1.37$, $p=.14$, respectively), suggesting that across-voxel patterns do not reliably discriminate between stimuli in pFs.

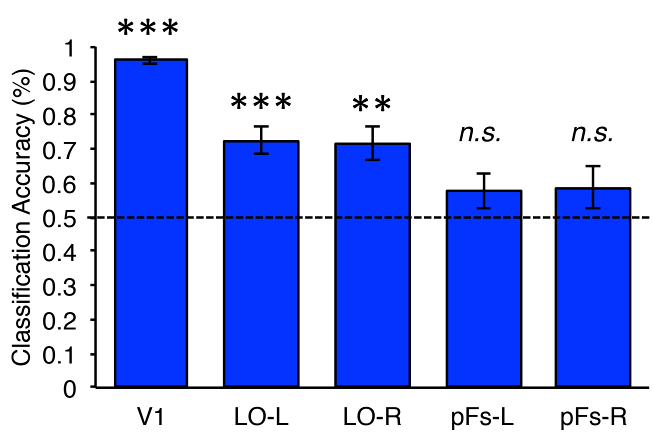

Figure S2. Average pairwise correlation-based classification (mean ± SEM) across stimuli within each ROI. V1 as well as LO in both hemispheres demonstrated above-chance classification, but this was not true for pFs in either hemisphere. $^{**}p < .01$; $^{***}p < .001$.

Supplementary Analysis 2: Consistency of between-stimulus similarity relations.

The previous analysis compared the patterns for the same stimulus across runs. We next asked whether the similarity relations between different stimuli were reliable across participants, by comparing the off-diagonal elements of each participant's MVP correlation matrix with one another. For each participant, we ran a single GLM on the entire dataset (all runs of the experiment) to estimate a multi-voxel pattern for each stimulus. T-values were used instead of beta values to optimize the stability of these estimates and enhance detection of across-participant similarities (Kravitz et al., 2010) (using beta values led to the same results). For each participant, we correlated the stimulus patterns with one another to generate a correlation matrix, and then correlated this correlation matrix (Pearson) with each other participant's correlation matrix. Averaging all of these between-participant correlations, we estimated a mean correlation representing the consistency of a given ROI's between-stimulus similarity structure across participants. To assess significance, we used permutation tests to compare this correlation to the distribution of 10,000 correlations computed when participants' correlation matrices were independently shuffled (we used a two-tailed test and a .01 criterion). The results of this analysis are show in Figure S3.

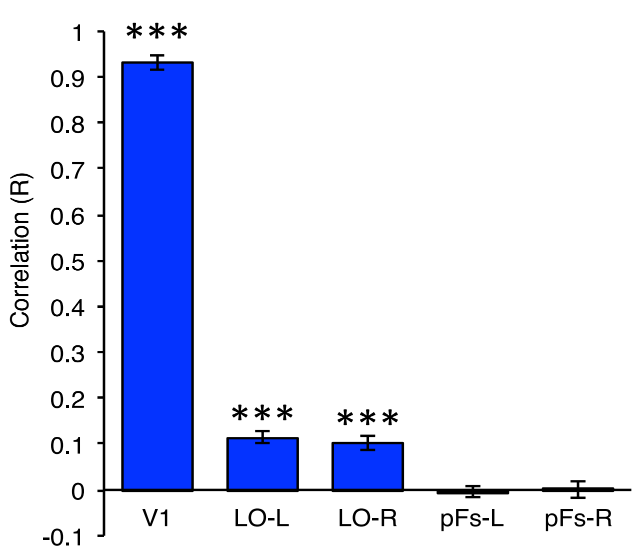

Figure S3. Consistency of between-stimulus similarity relations across participants. V1 as well as LO MVP correlation matrices were reliably correlated across participants, but pFs (in either hemisphere) was not. ***$p < .001$.

In V1, between-stimulus MVP correlations were extremely consistent across participants (R = 0.9308; 95% CI: [0.9014, 0.9602]). In LO, across-participant correlations were lower but reliably above what would be expected by chance (LO-L: R = 0.113; 95% CI: [0.087, 0.140]; LO-R, R = 0.101; 95% CI: [0.073, 0.131]). However, pFs did not show a reliable across-participant correlation (pFs-L: R = -0.003; 95% CI: [-0.028, 0.023]; pFs-R: R = 0.0011; 95% CI: [-0.035, 0.037]). This may be due to sensitivity of pFs representations to the idiosyncratic perceptual similarity judgments of individuals (Haushofer et al. 2008), or it may be a limitation of our scanning resolution in detecting structure at a finer scale. Given failure to detect reliable within-stimulus correlations and consistent between-stimulus similarity structure in pFs, all further analyses of multi-voxel patterns were restricted to LO and V1.

Supplementary Analysis 3: MVP-similarity searchlight analysis.

To determine whether any object-responsive regions outside of our ROIs were sensitive to the behavioral confusability or pixel overlap of our stimuli, we performed an MVPA searchlight analysis (Kriegeskorte, 2006; Rothlein & Rapp, 2014). In this procedure, we move a searchlight sphere (5-voxel diameter) throughout a predefined cortical search space, at each position generating an MVP-similarity matrix from the selected voxels and performing a regression-based MVP-similarity analysis to determine the relative contributions of behavioral confusability and pixel overlap in predicting the observed MVP-similarity patterns.

We follow the procedure described in detail in Rothlein & Rapp (2014). To define the search space, we used the object localizer data from each participant, normalized to talairach space and smoothed using a 4mm FWHM kernel, and entered it into a group random effects analysis GLM. We identified all "object-responsive voxels" as those voxels that showed significant activation to the presentations of objects (+Objects) at a threshold of $p<.01$ uncorrected with cluster threshold of 4 voxels. These voxels were used to defined a group-level mask that was reflected across the mid-saggital plane to create a bilaterally symmetric, large object-responsive search space that was used for all participants. This search space encompassed nearly the entirety of ventral visual cortex, extending up into inferior portions of occipito-parietal cortex.

For each participant, this search space was probed by a moving a searchlight volume (a sphere 5 functional voxels in diameter, encompassing a maximum of 125 voxels or minimum of 33 voxels, allowed for evaluating the edges of the search space) across the search space. At each position of the searchlight volume, the currently "highlighted" voxels were identified and the beta values for all 16 stimuli were extracted, z-scored, and correlated with one another to generate an observed MVP-similarity matrix. This observed MVP-similarity matrix was then entered into a regression where it was predicted as a function of the behavioral confusability and pixel overlap of stimuli (as well as a constant). The resulting beta values for 1) behavioral confusability and 2) pixel overlap were assigned to the voxel at the center of the searchlight, generating a beta map for each predictor. The searchlight was then moved by one voxel within the search space and this procedure was repeated until the entire search space had been analyzed.

Individual participant betas maps were then smoothed with a 2mm FWHM Gaussian kernel and then combined across participants to form a group map. To assess statistical significance at the group level, we performed a two-step procedure. First, we performed a group-level t-test comparing the betas for each predictor at each voxel to zero, producing a group t-map. This t-map was then corrected using an additional cluster size correction (voxelwise uncorrected: $p < 0.1$; cluster size corrected < 0.05) using BrainVoyager's Cluster Size Correction

Plugin. This produced a final group t-map of the voxels showing significant sensitivity to behavioral confusability and Pixel overlap, shown in the top and bottom panels of Fig S4, respectively.

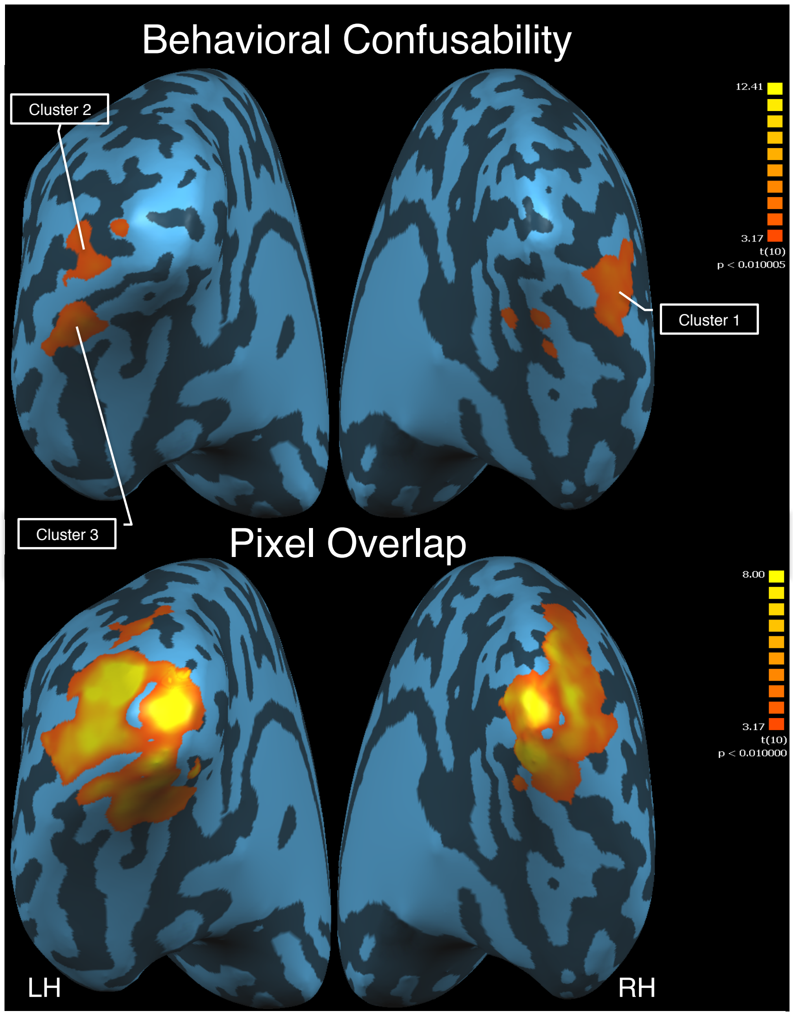

Figure S4. MVP-similarity searchlight analysis results. *Top panel.* Group-level t-map showing regions where MVP-similarity was significantly predicted by behavioral confusability. These 3 main clusters overlap with those defined as LO in our object localizer. *Bottom panel.* Regions where MVP-similarity was significantly predicted by pixel overlap. These regions include V1 but also extend to other early retinotopic regions.

The results of the MVPA searchlight analysis are consistent with those of our ROI analysis. The behavioral confusability t-map (Fig S4, top panel) revealed 3 major clusters, 1 in the RH and two in the LH, as well as some smaller regions in the RH visible in the figure.

Cluster 1 corresponded approximately to the LO-R (Cluster 1 talairach center of gravity ± 1 SD: 43.65 ± 3.55, -68.4 ± 5.72, -4.34 ± 6.19; across-participant average LO-R center of gravity 42.10, -70.40, -10.87]. Cluster 2 overlapped partially with LO-L, although its center of gravity was slightly more medial and posterior (Cluster 2 talairach center of gravity: -34.87 ± 3.73, -83.72 ± 3.96, -6.97± 3.39; across-participant average LO-L center of gravity: -43.35,-70.89,-8.30]. It should be noted that MVP correlations for the voxels in Cluster 2 were also significantly predicted by the pixel overlap model (Fig. S4, bottom panel). Cluster 3 corresponded more closely to LO-L and potentially a very posterior portion of pFs (Cluster 3 talairach center of gravity: -39.92 ± 3.65, -64.39 ± 4.47,-12.67 ± 3.72; across-participant average LO-L center of gravity: -43.35, -70.89, -8.30; across-participant average pFs-L center of gravity: -35.57, -46.68, -18.72]. No other regions showed significant sensitivity to behavioral confusability.

The pixel overlap t-map (Fig. S4, bottom panel) revealed 2 major clusters, encompassing V1 but also extending into adjacent early visual cortex. Although partially coextensive with cortex sensitive to behavioral confusability, sensitivity to pixel overlap was focused in the more posterior regions of cortex, and did not overlap at all (Cluster 1) or only slightly (Cluster 3) with 2 of the 3 main clusters identified in the behavioral confusability t-map.

In sum, given the exploratory nature and liberal thresholds applied in this analysis, we take the results to be consistent with our ROI analysis, suggesting that early visual regions are sensitive to the degree of pixel overlap between stimuli, whereas the more anterior, object-selective region LO is sensitive to the behavioral confusability of stimuli.

Supplementary Analysis 4: RS Computed over LO voxels contributing to MVP-similarity

In LO, we found that MVP-similarity is sensitive to the behavioral confusability of orientations, but RS is not. This dissociation was especially clear 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. However, it remains a possibility that at the individual voxel level, the voxels that respond most similarly to OPA reflections—thereby contributing most to the MVP-similarity results for OPA reflections—also tend to show RS.

To examine this possibility, we identified in each participant the LO voxels that responded more similarly to OPA reflections than to other non-identical stimulus pairs, and thereby contributed most to high MVP correlations for OPAs relative to other orientation relationships. We identified these voxels by computing an *OPA similarity ratio* for every voxel: OPA similarity ratio = (average absolute difference in response for all non-OPA pairs, excluding identity)/(average absolute difference in response for OPA pairs). Voxels with an OPA similarity ratio greater than 1 therefore responded more similarly (i.e., smaller differences in responses) to OPA reflections than other orientations. Across participants, 66.1% of voxels (range: 41% to 90%) showed OPA similarity ratios greater than 1.

We then asked whether voxels with an OPA similarity ratio greater than 1 showed RS for OPAs. When RS analyses were restricted to these voxels, we still did not observe RS to OPAs: the RS effect for voxels with OPA similarity ratio>1 did not differ from RS based on all voxels in LO, $t(10) = .16$, $p = .87$, and was not significantly different from zero, $t(10) = .38$, $p = .70$ (Fig. S5A).

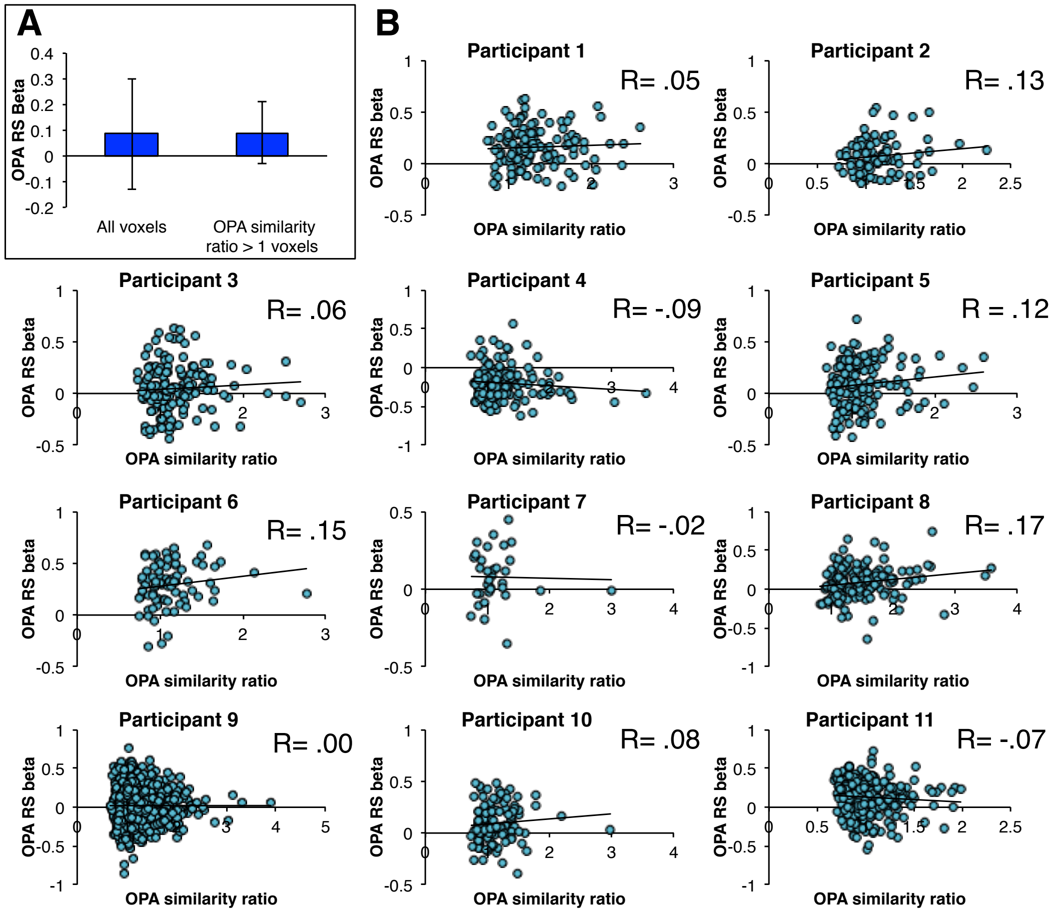

Figure S5. RS to OPAs is not related to similarity of voxel response to OPAs. (A) ROI-level RS for LO does not differ depending on whether all voxels (left) or only those voxels with OPA similarity ratio >1 (right) are included in the analysis. (B) By-voxel correlations between OPA similarity ratio (X-axis) and degree of RS to OPA reflections (Y-axis) for each participant. If increasing similarity of voxel response to OPA reflections predicted RS to OPAs, there should be a negative correlation between OPA similarity ratio and response to OPAs when presented in succession (OPA RS Beta). No such negative correlation was observed.

There may nonetheless be a correlation between similarity of response to OPAs and tendency to show RS for OPAs at the individual voxel level. If so, the OPA similarity ratio and the OPA beta from the RS analysis should be negatively correlated across voxels, with increasing similarity of response for OPA reflections (higher MVP-similarity) in a voxel being associated with decreasing response in that voxel when those orientation are presented in succession (greater RS effect). To test this, we calculated the correlation between the OPA similarity ratio and the OPA beta (representing the RS effect for repetition of OPAs) across all voxels in each participant's LO ROI. These correlations were then fisher-transformed and compared to zero using a two-tailed t-test. No relationship between RS and OPA similarity ratio was observed: the mean Pearson correlation was $R = .05$, which was not significant $t(10) = 1.97$, $p = .07$ **(**Fig. S5B**)**, indicating that the similarity of a voxel's response to OPA reflections does not predict its tendency to show RS for OPA reflections.

These results suggest that differences between MVP-similarity and RS are not simply due to averaging, and are consistent with other studies finding that RS and MVPA are not necessarily related at a voxel level (Sapountzis et al., 2010; Ward et al., 2013).

Supplementary Methods

Comparing OPA and EVA reflections: MVP-similarity.

We first identified the cells in the MVP correlation matrices that correspond to the multi-voxel correlations between Identical orientations, OPA reflections, and EVA reflections (Fig. S6A-1). For example, each OPA reflection cell contains the MVP correlation for a particular stimulus and its OPA reflection. In each participant's MVP correlation matrix, these cells are extracted and averaged (Fig. S6A-2). Averaging across participants, we generated a group average correlation for each orientation relationship (Fig. S6A-3). To assess the observed values for significance, we shuffled the three vectors of cells identifying each orientation relationship (Identical, OPA, and EVA) according to the same permutation (Fig. S6B-1) and re-computed the group average correlation (Fig. S6B-2). We repeat this process 10,000 times (Fig. S6B-3) to generate a distribution of group average correlations which may be compared to the values based on the true ordering to assess their significance (Fig. S6C-1).

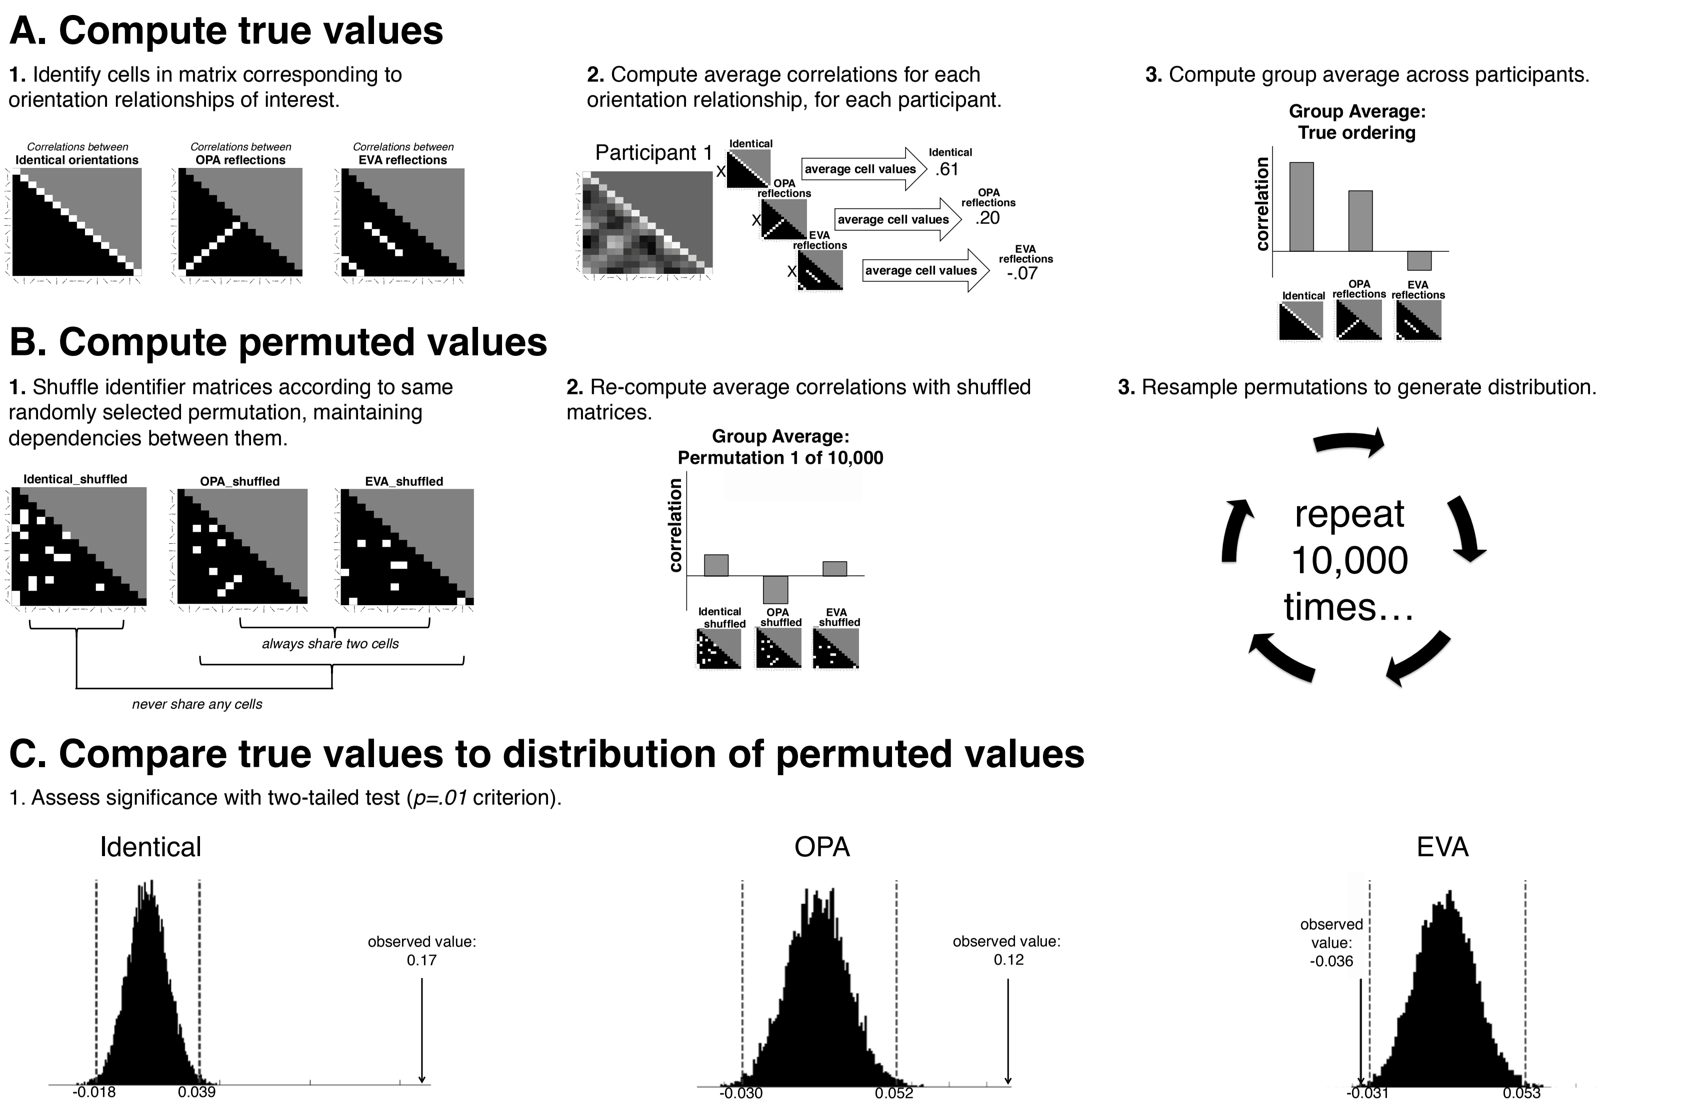

Figure S6. Steps of Comparing OPA and EVA reflections: MVP-similarity analysis. Dotted lines (C) mark the criterion for significance ($p = .01$, two-tailed) with corresponding correlations shown below.

Supplementary References

Aguirre, G. K. (2007). Continuous carry-over designs for fMRI. *NeuroImage*, *35*(4), 1480–1494.

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.

Haushofer, J., Livingstone, M. S., & Kanwisher, N. (2008). Multivariate patterns in object-selective cortex dissociate perceptual and physical shape similarity. *PLoS Biology*, *6*(7), e187.

Haxby, J. V. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*(5539), 2425–2430.

Kravitz, D. J., Kriegeskorte, N., & Baker, C. I. (2010). High-level visual object representations are constrained by position. *Cerebral Cortex*, *20*(12), 2916–2925.

Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(10), 3863–3868.

Rothlein, D., & Rapp, B. (2014). The similarity structure of distributed neural responses reveals the multiple representations of letters. *NeuroImage*, *89*, 331–344.

Sapountzis, P., Schluppeck, D., Bowtell, R., & Peirce, J. W. (2010). A comparison of fMRI adaptation and multivariate pattern classification analysis in visual cortex. *NeuroImage*, *49*(2), 1632–1640.

Ward, E. J., Chun, M. M., & Kuhl, B. A. (2013). Repetition suppression and multi-voxel pattern similarity differentially track implicit and explicit visual memory. *Journal of Neuroscience*, *33*(37), 14749–14757.