Research Article

Psychological Science
2016, Vol. 27(1) 3–11

DOI: 10.1177/0956797615598317
pss.sagepub.com

Neural Discriminability of Object Features Predicts Perceptual Organization

Emily J. Ward and Marvin M. Chun
Department of Psychology, Yale University

Abstract

How does the neural representation of simple visual features affect perceptual operations, such as perceptual grouping? If the strength of feature representations in the brain is indicative of how the perceptual system partitions information into visual elements, then identifying the underlying neural representation may determine why things look the way they do. During functional MRI, participants viewed objects that varied along three feature dimensions: shape, color, and orientation. Afterward, participants performed an independent perceptual-grouping task outside the scanner to measure the strength of feature grouping. In lateral occipital cortex, neural feature discriminability, characterized using functional MRI multivariate pattern classification, positively predicted feature grouping strength: The more distinct the neural representations of a particular feature, the stronger the grouping was for that feature outside the scanner. Thus, variation in neural feature representation can be quantified to predict perceptual organization.

Keywords

perception, representation, attention, fMRI

Received 1/9/15; Revision accepted 7/9/15

The perceptual system partitions and organizes the world into behaviorally relevant objects and groups. Many well-characterized Gestalt principles (Brooks, in press; Palmer, 1999; Wagemans, Elder, et al., 2012; Wertheimer, 1924/1950) account for some perceptual organization. But most visual objects in the environment comprise several features—for example, cars on the road come in different shapes, colors, and sizes. Quantifying how single features contribute to object formation—let alone quantifying the combined effect of multiple features—remains a perennially unsolved problem (Palmer & Beck, 2007; Wagemans, Elder, et al., 2012).

One way to understand how multiple features drive object formation is to measure how well the perceptual system parses certain features, then experimentally set those features against each other in a task designed to measure perceptual organization, such as feature grouping. However, because the goal of perception is to generate behaviorally relevant representations, different information will be selected and inhibited for different purposes (Carrasco, 2011; Chun, Golomb, & Turk-Browne, 2011). Therefore, any test of multifeature interaction must also consider the task relevance of the features, especially since simply attending to features can change the perception of their strength and appearance (Carrasco, Ling, & Read, 2004; Fuller & Carrasco, 2006).

To test whether the quality of perceptual feature representations is indicative of how the perceptual system partitions information into visual elements, we obtained a neural measure of feature representation and a perceptual measure of feature grouping. We first scanned participants while they viewed single objects that varied along three feature dimensions: shape, color, and orientation. We used multivariate pattern classification to characterize neural discriminability of object features in lateral occipital cortex (LOC), a high-level visual area sensitive to several Gestalt grouping phenomena, such as grouping of contour lines (Altmann, Bulthoff, & Kourtzi, 2003; Kourtzi, Tolias, Altmann, Augath, & Logothetis, 2003) and configural superiority (Kubilius, Wagemans, & Op de Beeck, 2011). After scanning, participants performed an

Corresponding Author:
Emily J. Ward, Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205
E-mail: emily.ward@yale.edu

independent behavioral task outside the scanner, the repetition-discrimination task (Palmer & Beck, 2007), designed to test how strongly different features perceptually group together when placed in competition with irrelevant features. The relative difference in performance when irrelevant features are compatible or incompatible with relevant features provides a measure of feature grouping strength (Palmer & Beck, 2007). By relating the measure of neural discriminability to the independent measure of feature grouping strength, we asked whether relevant object-feature representation could be quantified with functional MRI (fMRI) to predict basic perceptual operations, such as grouping.

We hypothesized that object features would become more neurally discriminable when they were relevant to participants in the scanner (Jehee, Brady, & Tong, 2011) and that features that were more neurally discriminable when relevant would group more strongly even in tasks performed outside the scanner. Specifically, high within-feature similarity and low between-features similarity should benefit neural classification *and* perception (Duncan & Humphreys, 1989), enabling fast detection of relevant features, as well as making the features more robust to conflicting feature information.

Method

Participants

Twenty-four paid participants (10 females, 14 males) gave written informed consent in compliance with guidelines of the Yale University Human Subjects Committee. This sample size was similar to that used in previous studies in our lab that link fMRI signals to behavior (e.g., Moore, Yi, & Chun, 2013; Ward, Chun, & Kuhl, 2013). Participants had normal or corrected-to-normal vision. Three additional participants were excluded from the study before data analysis (1 as a pilot subject, 1 for lacking subsequent behavioral data, and 1 for structural abnormalities).

Apparatus

We acquired fMRI data at the Magnetic Resonance Research Center at Yale University, on a 3T Siemens Trio scanner equipped with a standard 12-channel head coil. T1-weighted anatomical images were acquired using a 3-D magnetization-prepared rapid-acquisition gradient echo (MPRAGE) sequence—repetition time (TR) = 2,530 ms, echo time (TE) = 2.77 ms, time to inversion = 1,100 ms, voxel size = 1 × 1 × 1 mm, matrix size = 256 × 256 × 256. T2*-weighted images sensitive to blood-oxygen-level-dependent (BOLD) contrasts were acquired using a gradient-echo echo-planar pulse sequence (TR = 2,000 ms, TE = 25 ms, voxel size = 3.5 × 3.5 × 4 mm, matrix size = 64 × 64 × 34). Stimuli were presented using PsychoPy (Peirce, 2007) and displayed through an LCD projector on a rear-projection screen. Responses were recorded using a four-button fiber-optic response-pad system.

Stimuli

Stimuli were 27 images of striped geometrical items that varied by color (red, aqua, purple), shape (square, circle, triangle), and stripe orientation (120° diagonal, vertical, and 60° diagonal). To reduce low-level adaptation, we alternated the phase of the stripes every trial. The images measured 512 × 512 pixels. Images were presented at the center of the screen and subtended approximately 12° × 12° of visual angle.

Design and procedure

The experimental session included eight experimental scan runs (5 min 16 s each), followed by a behavioral perceptual-grouping test outside the scanner. In the scanner, participants were instructed before each run to perform one of four tasks while viewing the stimuli: (a) passive viewing, (b) indicating the color of the stimuli, (c) indicating the shape of the stimuli, or (d) indicating the orientation of the stripes of the stimuli. The task for the first two runs was always passive viewing to minimize attention to specific features (i.e., before participants were instructed to deploy specific feature-based attention). After the two passive-viewing scans, participants completed one scan for each task (color, shape, orientation) and then the second scan for each task in the reverse order (two scans of each task total). The task order was counterbalanced every 6 participants. Only scans in which participants performed an active task were used for the analysis.

Across all eight runs, participants saw 1,626 total trials. The sequence of trials was a predesigned Type 1 Index 1 27-element sequence based on Aguirre's (2007) continuous carryover design. This sequence is first-order counterbalanced (including repetition of elements) such that every item precedes (and follows) every other item. This permits measurement of the direct effect of each stimulus, independent of carryover effects from the other stimuli. The entire sequence included 28 repetitions of 27 items (plus one extra presentation of 1 item; 757 total), with 56 instances in which the screen was blank for 3,000 ms (to introduce jitter into the sequence). The sequence was repeated twice across the experiment. On experimental trials, each item was presented for 1,400 ms followed by a 100-ms blank screen (Fig. 1a). Participants indicated the task-relevant feature of each item by pressing a button-box key with their right index, middle, or ring finger.

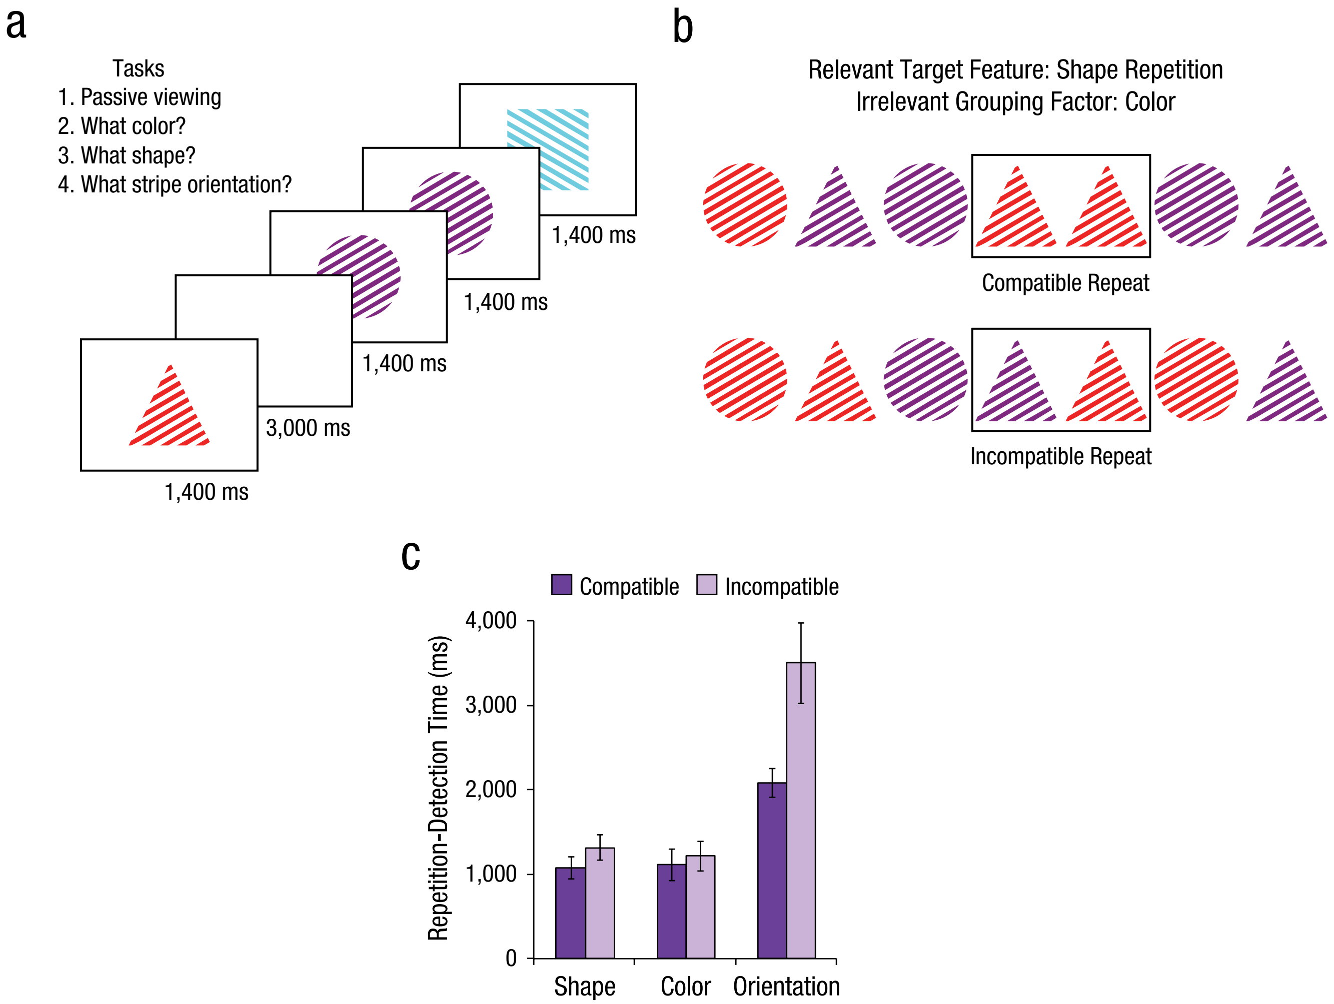

Fig. 1. Paradigm of the scanner tasks and paradigm and results for the repetition-detection task. In the scanner (a), participants were shown a sequence of items that differed in three dimensions (shape, color, and stripe orientation) while they performed one of four tasks: passive viewing, indicating the color of the items, indicating the shape of the items, or indicating the orientation of the stripes of the item. Each item was followed by a 100-ms blank screen (not shown here). Outside the scanner, participants performed the repetition-discrimination task (b), in which they were shown a series of nine items (seven are shown in this example) arranged in a row that alternated along one relevant feature dimension (in this example, shape). One of the items repeated consecutively, and participants' task was to indicate what was repeated. A second, but irrelevant feature alternated every two positions. On compatible trials, this feature was the same for both items in a target pair (thus facilitating identification), whereas on incompatible trials, this feature was different for each item. Reaction time on the repetition-detection task (c) is shown as a function of feature dimension and trial type. Error bars are 95% within-subjects confidence intervals.

After all eight scans were completed, participants performed a perceptual-grouping task outside the scanner using the same items. The repetition-discrimination task (Palmer & Beck, 2007) measures perceptual-grouping strength among features. Displays contained a single row of nine items drawn from the same stimuli used in the scanner tasks. These items alternated between two values in one feature dimension (e.g., circle and triangles in the dimension of shape), except for a consecutive pair in which the same feature value repeated (e.g., two triangles; Fig. 1b). Participants indicated what feature value repeated. However, in addition to the alternation between relevant features, a task-irrelevant feature dimension (such as color) served as an orthogonal grouping factor. On compatible trials, the target pair (in this example, two triangles) could have the same color (both red); thus, the orthogonal grouping factor (color) biased the target to be perceived as part of the same perceptual group (Fig. 1b, top row). On incompatible trials, the two items in the target pair could have different colors (purple and red); thus, the orthogonal grouping factor biased the target to be perceived as part of two different groups (Fig. 1b, bottom row).

The repetition-discrimination task had three blocks, containing 288 trials each (144 compatible trials and 144 incompatible trials), totaling 864 trials. For each block, participants detected a repetition of one of the three features (shape, color, or orientation). The trials were self-paced, but participants were encouraged to respond quickly without losing accuracy. Participants received auditory feedback for incorrect responses.

Data analysis

Cortical reconstruction of each participant's anatomical data was performed with FreeSurfer software (Fischl, 2012). Functional data processing was performed with AFNI (Cox, 1996) using the uber_subject.py processing stream. The first five volumes were removed from each scan to allow for scanner calibration. Functional images were corrected for motion (6 degrees of freedom) and for differences in slice timing. Volumes with more than 0.3 mm deviance in motion or more than 1% of the brain voxels were outliers in terms of intensity and therefore excluded. Data were aligned and registered to a standard Talairach space template (2-mm isotropic voxel size). Voxels were smoothed using a 4-mm full-width half-maximum Gaussian filter. For each run, a cubic polynomial was fit and removed to eliminate drift.

Our analyses were run on bilateral LOC regions of interest (ROIs) generated from the subject-specific cortical parcellations produced by FreeSurfer. As control regions, we also analyzed inferior parietal gyrus; response-inhibition ROIs, including anterior cingulate cortex and inferior frontal gyrus (Aron, Robbins, & Poldrack, 2004, 2014); and a V1 ROI taken from a probabilistic atlas (Wang, Mruczek, Arcaro, & Kastner, 2014; see Table S1 in the Supplemental Material). Data were modeled using a restricted maximum-likelihood (REML) regression implemented by AFNI to account for serial autocorrelation in time-series noise. Because of the high temporal pace of the stimulus presentation, accounting for autocorrelation was especially important. Each item in each task served as a separate regressor (27 items × 4 tasks = 108 total regressors in the model). Hemodynamic response function was fit to each stimulus presentation (which were treated as instantaneous) using a gamma basis function with its peak 4.7 s after stimulus onset (reflecting hemodynamic delay). This REML model resulted in beta values that were extracted for each regressor and for each voxel within the two bilateral LOC ROIs. These values composed the patterns of activity used in the subsequent analyses.

Pattern classification

We used multivariate pattern classification to characterize neural discriminability of object features. For each classification analysis, we trained a linear support vector machine (Pedregosa et al., 2011) using a leave-one-item-out cross-validation approach. An example workflow is as follows: To classify relevant versus irrelevant features, we sorted the patterns of beta values in the ROI on the basis of task (indicate shape, color, or orientation). There were 27 items for each task. We labeled the patterns for the items on the basis of the relevant feature (e.g., square, circle, or triangle for the shape task). The classifier was trained on 26 of these patterns and tested on the remaining pattern. This process repeated until all permutations of the set had been used for training and testing, at which point the average accuracy of feature classification was determined. Because there were three feature values for each feature dimension, chance accuracy was 33.33%. It is worth noting that each item always had two other features in common with the other items (e.g., a red square with vertical stripes shared a feature with all red items and with all items with vertical stripes, square or not), and the classifier must learn to ignore these similarities to successfully classify the item using just the currently relevant feature. As a result, our absolute classification-accuracy values are an underestimate of their representation in LOC (compared with accuracy values if the features were somehow presented in isolation).

Representational similarity

To ensure that our results generalized across different analytic approaches, we also analyzed our data using a pattern-similarity method (Kriegeskorte, Mur, & Bandettini, 2008; Ward et al., 2013). Similarity approaches are less sensitive but also do not rely on complex voxel-weight metrics and thus can be more interpretable. Similarity was computed using Pearson correlations between the patterns for each unique pair of the 27 items (351 values total). The resulting z-transformed r values correspond to either same feature value (e.g., two circles) or different feature values (e.g., a circle and a square or triangle), with the assumption that patterns for the same item are more similar than patterns for different items. Thus, the difference between same-value similarity and different-value similarity serves as an analogous measure of representational discriminability for that dimension.

Mixed-effects modeling

To relate the feature-classification results to the perceptual-grouping behavioral measure, we used mixed-effects analysis implemented in lme4 (Bates, Maechler, & Bolker, 2012) in R software (R Development Core Team, 2012). Mixed-effects modeling is a powerful statistical tool that offers many advantages over conventional t-test, regression, and variance and covariance analyses in sophisticated experimental designs. Our analysis approach with the mixed-effects models was to treat classification accuracy for relevant and irrelevant features as a predictor of behavior (perceptual-grouping strength). Linear mixed-effects modeling accommodated our design well, because we were able to model the relationship between classification accuracy and grouping strength for each combination of relevant and irrelevant features. To determine the significance of a factor (e.g., relevant-feature classification accuracy), we used the individual estimates for each factor (e.g., how well relevant-feature classification predicted grouping strength). Subject ($N = 24$) and feature combination (six total combinations) were included as random effects. Because the strength of the relationship between classification accuracy and perceptual grouping could vary by combination (i.e., it might be stronger for shape vs. orientation but less strong for shape vs. color), we included a random slope for classification accuracy, which also provides the regression estimates for each combination.

Results

Behavioral data

In the three active tasks performed in the scanner, participants were highly accurate at indicating the task-relevant feature of each item, with no significant differences among tasks in accuracy (color: $M = 94.33\%$, $SD = 1.3$; shape: $M = 93.99\%$, $SD = 1.7$; orientation: $M = 94.00\%$, $SD = 2.4$), $F(2, 46) = 0.032$, $p = .969$, $\eta_p^2 = .001$, and reaction time (RT; color: $M = 626$ ms, $SD = 97$; shape: $M = 622$ ms, $SD = 85$; orientation: $M = 618$ ms, $SD = 99$), $F(2, 46) = 0.31$, $p = .7348$, $\eta_p^2 = .01$. A strong priming effect occurred when the relevant feature repeated (repeat: $M = 564$ ms, $SD = 80$; no repeat: $M = 649$ ms, $SD = 107$), $F(1, 23) = 76.53$, $p < .001$, $\eta_p^2 = .77$. Furthermore, a strong priming effect occurred when irrelevant features repeated, even when there was no repetition of the task-relevant feature (no irrelevant repetition: $M = 655$ ms, $SD = 110$; one irrelevant repetition: $M = 646$ ms, $SD = 109$; two irrelevant repetitions: $M = 647$ ms, $SD = 106$), $F(1, 23) = 4.38$, $p = .048$, $\eta_p^2 = .16$.

fMRI classification

In bilateral LOC, overall classification accuracy for all features (39.65%, $SD = 7.59$) was greater than chance (33.33%), $t(23) = 4.08$, $p < .001$, $d = 0.83$. There was a main effect of feature on classification accuracy, $F(2, 46) = 12.84$, $p < .001$, $\eta_p^2 = .36$. Shape (44.19%, $SD = 9.97$), $t(23) = 5.34$, $p < .001$, $d = 1.09$, and color (39.60%, $SD = 10.26$), $t(23) = 2.99$, $p = .006$, $d = 0.61$, could be classified above chance, but orientation could not (35.16%, $SD = 7.00$), $t(23) = 1.28$, $p = .21$, $d = 0.26$. There was no interaction between feature relevance and dimension, $F(2, 46) = 1.21$, $p = .307$, $\eta_p^2 = .05$ (Fig. 2b): Within each task, task-relevant features (e.g., circle, square, or triangle during the shape task scans) were better classified than task-irrelevant features (relevant: $M = 42.21\%$, $SD = 12.20$; irrelevant: $M = 38.80\%$, $SD = 9.87$), $F(1, 23) = 12.27$, $p = .002$, $\eta_p^2 = .35$ (Fig. 2a), consistent with previous research demonstrating task-relevant feature enhancement (Jehee et al., 2011).

Representational similarity analyses also revealed that patterns were discriminative. Specifically, patterns for the same feature value (e.g., all circles) were more correlated than patterns for different feature values (same: $r = .50$, $SD = .24$; different: $r = .49$, $SD = .23$), $F(1, 23) = 14.81$, $p < .001$, $\eta_p^2 = .39$, except in the orientation dimension, $t(23) = 0.16$, $p = .874$, $d = 0.03$. In general, representational similarity (r for same values − r for different values) was highly correlated with classifier accuracy ($r = .85$, $p < .001$) and yielded similar results, although the difference in similarity between relevant and irrelevant features was only marginally significant, $F(1, 23) = 3.65$, $p = .069$, $\eta_p^2 = .14$.

Perceptual grouping (repetition discrimination)

In the postscan perceptual-grouping behavioral task, overall repetition-detection accuracy was high (92.20%, $SD = 10.97$). Accuracy was higher when relevant-feature repetitions were grouped by the irrelevant feature (compatible trials: $M = 92.76$, $SD = 10.09$) than when relevant-feature repetitions were not grouped by the irrelevant feature (incompatible trials: $M = 91.64$, $SD = 11.83$), $F(1, 23) = 8.05$, $p = .009$, $\eta_p^2 = .26$. Similarly, correct RTs were faster for repetitions on compatible trials ($M = 1{,}405$ ms, $SD = 657$) than for repetitions on incompatible trials ($M = 1{,}985$ ms, $SD = 1{,}358$), $F(1, 23) = 86.84$, $p < .001$, $\eta_p^2 = .79$. However, there was a significant interaction between grouping type and task, $F(2, 46) = 41.96$, $p < .001$, $\eta_p^2 = .65$ (Fig. 1c), which indicated that the difference in RTs between compatible and incompatible trials was much greater for orientation ($M = 1{,}399$ ms, $SD = 862$) than for either color ($M = 101$ ms, $SD = 223$), $t(23) = 6.98$, $p < .001$, $d = 1.46$, or shape ($M = 239$ ms, $SD = 178$), $t(23) = 6.16$, $p < .001$, $d = 1.28$. To minimize the differences between compatible and incompatible trials among the three features, we defined grouping strength as the (log-transformed) ratio of compatible-trial RTs to incompatible-trial RTs. Because the log transformation yields negative values for stronger grouping (i.e., faster compatible-trial RT and slower incompatible-trial RT), we took the additive inverse of the values to make the measure more intuitive.

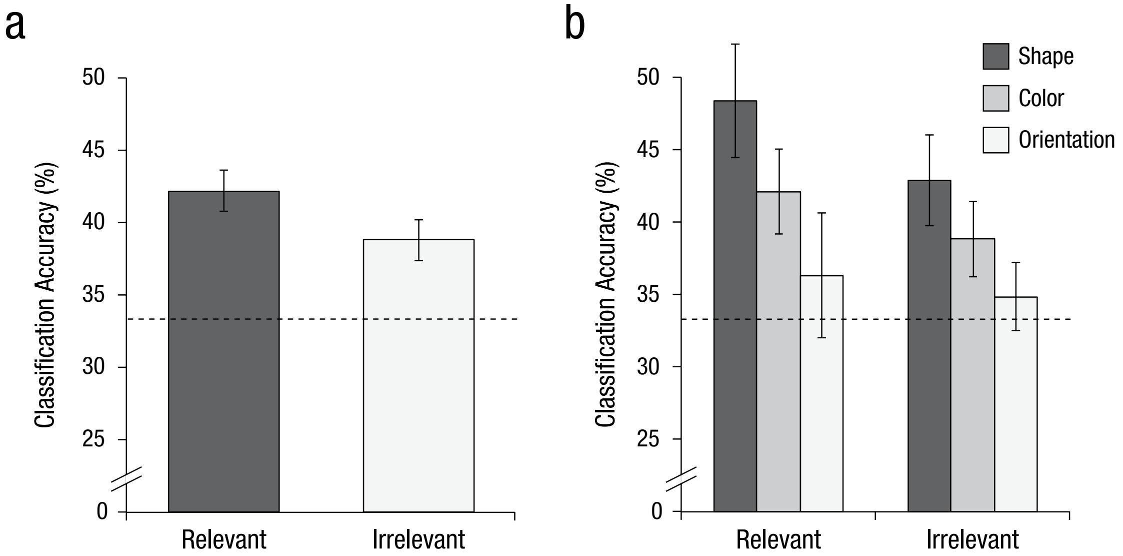

Fig. 2. Classification accuracy. Results are shown for task-relevant and task-irrelevant features (a) across feature dimensions and (b) separately for each feature dimension. The dashed lines indicate chance-level performance. Error bars are 95% within-subjects confidence intervals.

Predicting grouping strength from fMRI classification accuracy

To test whether task-relevant visual representations influenced perceptual grouping, we related relevant-feature multivariate classification (as our neural measure of discriminability) in the scanner to grouping strength outside the scanner (as our behavioral measure of perceptual grouping) for each subject. This was computed for each of the six combinations of relevant and irrelevant features: three relevant features (color, shape, orientation) × 2 possible irrelevant features (e.g., shape and orientation when the relevant feature was color). In this example, relevant classification accuracy for color and irrelevant classification accuracy for shape were related to the grouping-strength value for color versus shape. The same classification value for color would also be paired with irrelevant classification accuracy for orientation and related to the grouping strength for color versus orientation. If the characteristics of feature representation that benefit neural classification also underlie perceptual grouping, greater classification accuracy for the relevant feature than for an irrelevant feature should lead to stronger perceptual grouping of the relevant feature. We used a mixed-effects analysis to accommodate the different combinations, treating relevant and irrelevant classification accuracy as a predictor of perceptual-grouping strength.

We found a clear positive relationship between relevant-feature classification accuracy and grouping strength (Fig. 3). Across the combinations, classification accuracy for the relevant feature was a significant predictor of grouping strength ($b = 0.313$, $t = 2.56$, $p = .0320$, 95% confidence interval, or CI = [0.07, 0.55]). Classification accuracy for the irrelevant feature did not predict grouping strength ($b = -0.09$, $t = 0.87$, $p = .410$, 95% CI = [−0.30, 0.12]).

Furthermore, representational similarity (*r* for same values − *r* for different values) for relevant features was a marginal predictor of grouping strength ($b = 1.25$, $t = 1.97$, $p = .051$, 95% CI = [0.01, 2.49]). Representational similarity for the irrelevant feature did not predict grouping strength, although it trended in the opposite direction ($b = -1.16$, $t = 1.81$, $p = .07$, 95% CI = [−2.41, 0.09]).

To rule out the possibility that this relationship was simply driven by observers' ability to consciously ignore the features, we related RTs from the scanner tasks (when observers were indicating the relevant feature value for single items) to grouping strength. Naturally, observers' mean RT on the two tasks was correlated ($r = .62$), $t(22) = 3.677$, $p = .001$, but critically, we found no correlation between mean RT from the scanner tasks and *grouping strength*, our measure of interest ($r = .26$), $t(22) = 1.24$, $p = .23$. We also analyzed a more specific measure of irrelevant-feature processing from the scanner: irrelevant-feature priming—that is, how an observers' response to the relevant feature was affected by his or her processing of irrelevant information (RT when no irrelevant information repeated − RT when any irrelevant information repeated). We found no correlation between this measure and grouping strength ($r = -.19$), $t(22) = 0.93$, $p = .36$. In fact, numerically, this correlation went in the opposite direction: As noted earlier, irrelevant information facilitated responding in the single-item tasks, while

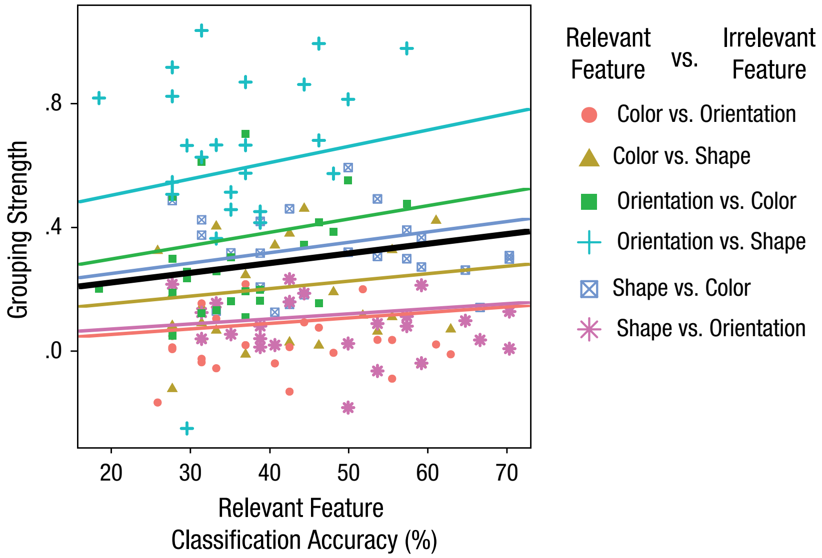

Fig. 3. Scatterplot showing the relationship between classification accuracy for relevant features and grouping strength for each of the six possible comparisons of relevant versus irrelevant features. Grouping strength was defined as the log-transformed ratio of compatible-trial reaction times to incompatible-trial reaction times. Each colored line represents the regression estimates for a particular combination of features, and the black line corresponds to the regression estimates for the overall relationship between the two measures.

it helped or hurt differentially in the grouping task, depending on features. Thus, the positive relationship between relevant-feature classification accuracy and grouping strength does not likely reflect just a general ability to ignore the same features between tasks. Instead, it shows that the more perceptually discriminable the relevant features are in object-selective cortex, the stronger the features group together in a perceptual-grouping task outside the scanner.

General Discussion

We quantified differences in perceptual feature discriminability using fMRI and showed that these measures predicted perceptual grouping in a separate behavioral task. Feature discriminability could be measured by how well we could classify features on the basis of their multivoxel patterns in LOC. Pattern-classification accuracy of features was higher for task-relevant features than for irrelevant features, consistent with prior research showing that task-relevance improves encoding for relevant features (Jehee et al., 2011). Notably, variation in pattern-classification accuracy for relevant features, measured for single objects viewed in the scanner, predicted how strongly these features grouped together perceptually, as assessed outside the scanner with a completely independent task widely used to measure perceptual grouping (Palmer & Beck, 2007; Vickery, 2008; Vickery & Jiang, 2009). Specifically, the more discriminable the features were in the brain, the stronger the features grouped together in a behavioral perceptual-grouping task. This indicates that the neural representations that drive successful pattern classification affect perception in general.

How might these neural representations determine what people see? Consider cases of high feature discriminability and low feature discriminability: A highly discriminable feature dimension would have strong within-feature grouping (e.g., squares would be perceived as highly similar to each other, and thus a repetition of two squares would be easy to see) and strong between-feature separation (e.g., squares would be perceived as very different from circles). In the domain of perceptual grouping, compatible irrelevant information would facilitate repetition detection (e.g., a repetition of squares will be aided by a compatible color grouping), but because the repetitions (within feature) are strongly grouped to begin with, incompatible irrelevant information would not drastically impair performance. This would lead to strong grouping in this feature dimension.

In the other case, a less discriminable feature dimension would have weak within-feature grouping (e.g., squares may not be as readily perceived as similar to each other, and thus a repetition of two squares would be more difficult to see) and weak between-feature separation (e.g., squares would be perceived as less different from circles). In the domain of perceptual grouping, a repetition of features would be difficult to detect, whether the irrelevant information is compatible or incompatible. This would lead to weak grouping in this feature dimension.

An alternative possibility is that these results reflect different abilities to inhibit responses to irrelevant information, rather than differences in perceptual processing. We distinguish between two types of inhibition described by Aron and colleagues (2004, 2014). At the perceptual level, attentional weighting—enhancement of task-relevant features and inhibition of task-irrelevant features—determines the perceptual discriminability of individual features; we certainly view our fMRI results as measuring this level of perceptual representation. However, behavior is also affected by response inhibition, which occurs postperceptually and is best demonstrated in tasks such as the Wisconsin Card Sorting Test or the Stroop task. This interpretation predicts that perceptual processing of task-relevant and task-irrelevant features should be similar across tasks because inhibition is occurring postperceptually.

Three aspects of our results favor a perceptual-level explanation rather than a response-inhibition account. First, our effects are specific to object-selective cortex and were not observed in ROIs such as right anterior cingulate cortex and right inferior frontal gyrus, which are implicated in cognitive control (Aron et al., 2004, 2014), or even ROIs that guide attention, such as inferior parietal cortex (see Table S1). Second, the response-inhibition account predicts that task-irrelevant features should slow down responses, but we observed the opposite result. In the single-item task in the scanner, irrelevant-feature information facilitated response time, which reveals perceptual priming (Chun & Nakayama, 2000; Kristjánsson & Campana, 2010). In the grouping task, whether irrelevant information helped or hurt depended on the specific feature combination and on whether or not the irrelevant information was compatible with the repetition. Finally, although the response-inhibition account specifically predicts that observers' responses in the tasks will be correlated because the tasks both require inhibition of the same irrelevant information, we found no relationship between irrelevant-feature priming and grouping strength across our two tasks. Thus, rather than response inhibition, our data suggest that feature discriminability and behavior are dynamically affected by task relevance at the perceptual level.

Thus, our findings support a dynamical systems approach to perceptual organization (see Wagemans, Feldman, et al., 2012, for a review), in which perception is not simply a process of going from static input to static output but depends on a combination of such factors as the physical salience of the feature, current goals, and selection history. We have shown that fMRI can quantify this dynamic variance in neural feature representation across individuals and can predict basic perceptual operations, such as grouping, in an independent task.

Author Contributions

E. J. Ward developed the study concept. E. J. Ward and M. M. Chun both contributed to the study design. E. J. Ward collected and analyzed the data. Both authors wrote the final version of the manuscript and approved it for submission.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

Funding

This work was supported by the Yale University Faculty of Arts and Sciences MRI Program funded by the Office of the Provost and the Department of Psychology.

Supplemental Material

Additional supporting information can be found at http://pss.sagepub.com/content/by/supplemental-data

References

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

Altmann, C. F., Bulthoff, H. H., & Kourtzi, Z. (2003). Perceptual organization of local elements into global shapes in the human visual cortex. *Current Biology*, *13*, 342–349.

Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, *8*, 170–177.

Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2014). Inhibition and the right inferior frontal cortex: One decade on. *Trends in Cognitive Sciences*, *18*, 177–185.

Bates, D., Maechler, M., & Bolker, B. (2012). lme4: Linear mixed-effects models using S4 classes (R package version 0.999999 – 0). Retrieved from http://www.inside-r.org/packages/lme4

Brooks, J. L. (in press). Traditional and new principles of perceptual grouping. In J. Wagemans (Ed.), *Oxford handbook of perceptual organization*. Oxford, England: Oxford University Press.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*, 1484–1525.

Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience*, *7*, 308–313.

Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology*, *62*, 73–101.

Chun, M. M., & Nakayama, K. (2000). On the functional role of implicit visual memory for the adaptive deployment of attention across views. *Visual Cognition*, *7*, 65–81.

Cox, R. W. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, *29*, 162–173.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*, 433–458.

Fischl, B. (2012). FreeSurfer. *NeuroImage*, *62*, 774–781.

Fuller, S., & Carrasco, M. (2006). Exogenous attention and color perception: Performance and appearance of saturation and hue. *Vision Research*, *46*, 4032–4047.

Jehee, J. F., Brady, D. K., & Tong, F. (2011). Attention improves encoding of task-relevant features in the human visual cortex. *Journal of Neuroscience*, *31*, 8210–8219.

Kourtzi, Z., Tolias, A. S., Altmann, C. F., Augath, M., & Logothetis, N. K. (2003). Integration of local features into global shapes: Monkey and human fMRI studies. *Neuron*, *37*, 333–346.

Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis—Connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *2*, Article 4. Retrieved from http://journal.frontiersin.org/article/10.3389/neuro.06.004.2008/full

Kristjánsson, Á., & Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search tasks. *Attention, Perception, & Psychophysics*, *72*, 5–18. doi:10.3758/APP.72.1.5

Kubilius, J., Wagemans, J., & Op de Beeck, H. (2011). Emergence of perceptual Gestalts in the human visual cortex: The case of the configural-superiority effect. *Psychological Science*, *22*, 1296–1303.

Moore, K., Yi, D. J., & Chun, M. M. (2013). The effect of attention on repetition suppression and multivoxel pattern similarity. *Journal of Cognitive Neuroscience*, *25*, 1305–1314.

Palmer, S. E. (1999). *Vision science: Photons to phenomenology*. Cambridge, MA: MIT Press.

Palmer, S. E., & Beck, D. (2007). The repetition discrimination task: An objective method for studying perceptual grouping. *Perception and Psychophysics*, *69*, 68–78.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Olivier, G., . . . Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830.

Peirce, J. W. (2007). PsychoPy—Psychophysics software in Python. *Journal of Neuroscience Methods*, *162*, 8–13.

R Development Core Team. (2012). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing.

Vickery, T. J. (2008). Induced perceptual grouping. *Psychological Science*, *19*, 693–701.

Vickery, T. J., & Jiang, Y. V. (2009). Associative grouping: Perceptual grouping of shapes by association. *Attention, Perception, & Psychophysics*, *71*, 896–909.

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & von der Heydt, R. (2012). A century of Gestalt psychology in visual perception: I. Perceptual grouping and figure-ground organization. *Psychological Bulletin*, *138*, 1172–1217.

Wagemans, J., Feldman, J., Gepshtein, S., Kimchi, R., Pomerantz, J. R., Van der Helm, P. A., & Van Leeuwen, C. (2012). A century of Gestalt psychology in visual perception: II. Conceptual and theoretical foundations. *Psychological Bulletin*, *138*, 1218–1252.

Wang, L., Mruczek, E. B. R., Arcaro, M. J., & Kastner, S. (2014). Probabilistic maps of visual topography in human cortex. *Cerebral Cortex*. Advance online publication. doi:10.1093/cercor/bhu277

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

Wertheimer, M. (1950). Gestalt theory. In W. D. Ellis (Ed.), *A source book of Gestalt psychology* (pp. 1–11). New York, NY: Humanities Press. (Original work published 1924)