

Category Learning Increases Discriminability of Relevant Object Dimensions in Visual Cortex

Jonathan R. Folstein, Thomas J. Palmeri and Isabel Gauthier

Department of Psychology, Vanderbilt University, Nashville, TN 37240, USA

Address correspondence to Jonathan R. Folstein. Email: jonathan.r.folstein@gmail.com.

Learning to categorize objects can transform how they are perceived, causing relevant perceptual dimensions predictive of object category to become enhanced. For example, an expert mycologist might become attuned to species-specific patterns of spacing between mushroom gills but learn to ignore cap textures attributable to varying environmental conditions. These selective changes in perception can persist beyond the act of categorizing objects and influence our ability to discriminate between them. Using functional magnetic resonance imaging adaptation, we demonstrate that such category-specific perceptual enhancements are associated with changes in the neural discriminability of object representations in visual cortex. Regions within the anterior fusiform gyrus became more sensitive to small variations in shape that were relevant during prior category learning. In addition, extrastriate occipital areas showed heightened sensitivity to small variations in shape that spanned the category boundary. Visual representations in cortex, just like our perception, are sensitive to an object's history of categorization.

Keywords: acquired distinctiveness, category learning, fMRI, object representation, perceptual learning

Introduction

Objects differ from one another along multiple perceptual dimensions. These dimensions can be simple, like size and brightness (Goldstone 1994), more complex but local, like object parts (Nosofsky 1986; Sigala and Logothetis 2002), or more global properties of object shape (Goldstone and Steyvers 2001; Freedman et al. 2003; Jiang et al. 2007; Gureckis and Goldstone 2008). Learning that objects belong to different categories can cause them to be perceived differently and many models account for category learning by stretching this multidimensional space along relevant dimensions (Nosofsky 1986; Kruschke 1992). This stretching causes objects that differ along a relevant dimension to be less similar to one another.

Beyond facilitating categorization, this dimensional modulation of object representations can have lasting consequences on how we perceive objects: Objects that vary along a dimension relevant to previously learned categories can remain more discriminable in perceptual tasks that do not require categorization, a phenomenon called “acquired distinctiveness.” For example, Goldstone found that subjects trained to categorize shaded squares according to either size or brightness were afterward better able to visually discriminate them along whichever dimension had been relevant during category learning (Goldstone 1994). This increase in discriminability was observed globally for all values along the relevant dimension but was especially large for local regions of the stimulus space around the category boundary—that is, for similar pairs of

stimuli that belonged to different categories. (In this and other studies, Goldstone and colleagues have also found effects of acquired equivalence, or decreases in discriminability along the irrelevant dimension. Interestingly, we have found little evidence for acquired equivalence in our behavioral studies, so we focus in this paper on acquired distinctiveness.)

Enhanced visual discriminability along relevant dimensions following category learning has also been observed for 2D shape spaces (Goldstone and Steyvers 2001; Op de Beeck et al. 2003; Hockema et al. 2005; Notman et al. 2005; Gureckis and Goldstone 2008; Folstein et al. Forthcoming). In some cases, these shape spaces were defined in terms of relatively basic shape dimensions that likely had psychological (and neural) representations prior to any category learning (Op de Beeck et al. 2003). In other cases where shape spaces were created by morphing complex objects, the dimensions that define those spaces are far less clear and may not have existed prior to category learning (Goldstone and Steyvers 2001; Gureckis and Goldstone 2008; Folstein et al. Forthcoming; see also Notman et al. 2005). In either case, objects that differ along dimensions relevant to the learned categories are more perceptually discriminable after learning, suggesting that category learning somehow enhances, and may even create, representations of those relevant object dimensions. Where in the brain do these representational changes take place?

Objects are thought to be represented in the ventral stream of the visual system (Grill-Spector 2003). It therefore seems possible that category-specific enhancements to object perception as a result of category learning might be accompanied by changes to object representations in the ventral stream. On the other hand, it is also possible that these observed behavioral changes in discriminability are not caused by changes in visual object representations per se but emerge from neural representations that are far more abstract. For example, semantic representations in the superior temporal gyrus or prefrontal cortex or episodic memory representations in the hippocampus could mediate acquired distinctiveness observed behaviorally.

Evidence for acquired distinctiveness in the ventral stream is currently mixed. Some neurophysiology and functional magnetic resonance imaging (fMRI) studies have suggested that relevant dimensions can be emphasized in visual cortex. After monkeys learned to categorize multidimensional objects, IT neurons were more sensitive to variations along a relevant dimension than along an irrelevant dimension (Sigala and Logothetis 2002; De Baene et al. 2008). In a study using fMRI, human subjects learned to categorize objects varying in shape and motion according to rules in which either shape or motion were relevant (Li et al. 2007). After training, fMRI pattern classification was used to show that information about stimulus category was enhanced in brain areas representing shape or motion when subjects categorized according to shape or

motion, respectively. In another fMRI study, Reber et al. (2003) observed qualitatively different patterns of fMRI adaptation in visual cortex for category members and nonmembers following either implicit or explicit learning of dot pattern categories, suggesting qualitatively different systems for learning and representing perceptual categories (but see Gureckis et al. 2011; Nosofsky et al. 2012).

Importantly, the studies reviewed so far focused on brain activity measured while subjects actively categorized objects. Attention to task-relevant features is known to have strong effects on the visual system (e.g., Corbetta et al. 1991). Thus, it is possible that these effects were due to some flexible top-down modulation of visual cortex present only during an active categorization task and not on a more stable change of object representations detectable even when attention is not directed toward category-relevant features. To understand whether changes in the visual system underlie the acquired distinctiveness observed in behavioral studies, it is also important to show that these neural effects can be observed in a task where the learned categorization is irrelevant, in the sense that dimensions relevant for categorization are irrelevant for the task at hand (Gauthier and Palmeri 2002).

Studies that have explicitly examined the neural consequences of category learning outside of active object categorization have found no selective enhancement of relevant dimensions in visual cortex (Jiang et al. 2007; Gillebert et al. 2008; van der Linden et al. 2010). For example, Jiang et al. (2007) used fMRI adaptation to measure changes in neural representations after category learning. Jiang et al. trained subjects to categorize complex objects residing in a morph space. They measured fMRI adaptation to pairs of objects before training and after training while subjects engaged in a task requiring no categorization, where they simply had to judge the location of an object on the screen. There was as much release from adaptation for sequentially presented pairs of objects in different categories as for pairs of objects within the same category, suggesting no enhancement in neural representations around the category boundary. Using analogous methods, others also failed to find evidence for systematic changes in neural discriminability in object-sensitive regions of visual cortex as a result of category learning (Gillebert et al. 2008; van der Linden et al. 2010).

In sum, the evidence for a neural locus of acquired distinctiveness within the visual system seems weak. As a result, there appears to be growing consensus that category learning has little impact on the visual system. While there is some evidence of modulation within visual cortex while objects are being actively categorized (Sigala and Logothetis 2002; De Baene et al. 2008), even this is not always obtained. Jiang et al. (2007) did not find effects of categorization in visual cortex when subjects were actively categorizing (see also Freedman et al. 2003). These results and others are consistent with a view that object category learning is largely nonvisual (Freedman and Miller 2008; Seger and Miller 2010), in that the ventral stream represents a vocabulary of shapes irrespective of category that is fed forward to more flexible areas capable of adaptively representing object categories (Serre et al. 2007; Cromer et al. 2010; Roy et al. 2010). By this account, while shape representations in visual cortex could develop sharper tuning with experience, these experience-dependent changes are insensitive to the type of experience, such as whether particular object dimensions were relevant or irrelevant to

previously learned categories. Here, we revisit this general conclusion and provide evidence that visual representations of complex objects in visual cortex can show acquired distinctiveness as a result of category learning when objects are not being actively categorized. While frontal and parietal areas can clearly contribute to category learning, our results demonstrate much more robust effects within visual areas than prior work.

One key feature of our approach is that first, we provide behavioral evidence for acquired distinctiveness before testing for its neural correlates in the scanner. In some previous cases (e.g., Gillebert et al. 2008; van der Linden et al. 2010) behavioral tests were simply not done. In other cases, no behavioral evidence for acquired distinctiveness was observed, yet these studies went on to look for (unsuccessfully) neural evidence for acquired distinctiveness anyway (Jiang et al. 2007). We should note that all of these studies (see also Freedman et al. 2003) used morphing methods that differed in subtle yet important ways from prior behavioral work, including some of our own, that has obtained behavioral evidence for acquired distinctiveness (Goldstone and Steyvers 2001; Folstein et al. Forthcoming). (While the differences between these morphing methods are technical and beyond the scope of this paper, we very briefly summarize them here: In a nutshell, the morphing method used by Jiang et al. (2007) and several others [Freedman et al. 2003; Gillebert et al. 2008; van der Linden et al. 2010] makes use of the full space of possible blends of 4 morphparents, likely making all shape variance relevant. The method used by Goldstone and colleagues [Goldstone and Steyvers 2001; Gureckis and Goldstone 2008] factorially blends between 2 morphlines, each created by blending 2 of the 4 parents. The factorial method allows the categorizer to attend to some shape variance while ignoring other shape variance, resulting in increased sensitivity to relevant dimensions as well as the ability to create novel dimensions [for further details, see Folstein et al. Forthcoming].) We argue that a prerequisite for finding any neural locus of acquired distinctiveness is first obtaining behavioral evidence for acquired distinctiveness. That means using experimental procedures and stimuli that produce the desired behavior.

After confirming acquired distinctiveness behaviorally for our stimulus set, we conducted an fMRI adaptation study to test if these psychophysically measured perceptual enhancements were accompanied by neural enhancements in representations for relevant dimensions in the ventral stream of the visual system. fMRI adaptation can be used to measure the degree of similarity between neural representations. When a population of neurons is stimulated twice, for instance by presenting the identical visual stimulus twice, the population of neurons will fire less the second time (Sawamura et al. 2006). In the broad strokes, this relationship seems to be preserved in fMRI BOLD activation as well: When two identical visual stimuli are presented twice, BOLD activation is less the second time. This is the phenomenon of fMRI adaptation or repetition suppression (Grill-Spector et al. 2006). When two stimuli are completely different, the amount of BOLD activation elicited by the second stimulus will be relatively unaffected by the first stimulus, a phenomenon known as “release from repetition suppression.” The more similar two visual stimuli are, the more repetition suppression to the second stimulus is seen in the visual system (Panis et al. 2008; Drucker et al. 2009). Our hypothesis is that, after category learning, visual representations of stimuli that differ along relevant dimensions will be less

similar than visual representations of stimuli that differ along irrelevant dimensions. Therefore, stimulus pairs that differ along relevant dimensions should show less repetition suppression than stimulus pairs that differ along irrelevant dimensions.

Materials and Methods

Subjects

Subjects were screened for how well they learned the categories (>80% accuracy in category learning, based on the reasonable intuition that greater categorization accuracy would lead to greater perceptual learning); a total of nine subjects did not pass the screening. Data were obtained from 20 subjects (11 females, average 25 years old), who took part in the fMRI experiment for financial compensation. An additional six subjects passed the screening and participated but were dropped due to poor performance within the scanner (<55% accuracy on at least 2 runs of either the location task or the category task [Performance below 55% accuracy was judged to be at chance.]), requesting to leave the scanner due to fatigue, or loss of 2 or more runs due to excessive motion (greater than 4.5 mm) or ghosting.

Materials

To begin with, multidimensional scaling (MDS) was used to select 4 car images, which we refer to as the parent images A-D, that had approximately equal similarity to one another from among 30 total images (selected from a collection of 3D computer models available online at www.doschdesign.com/products/3d/Lo-Poly_Cars_V1-2.html). These 4 parents (Fig. 1) were further processed using Adobe Photoshop to remove unnecessary details that were interfering with smooth morphing between parents (mostly limited to changing glossy reflective surfaces into matte surfaces). Although it is possible that these small modifications to the cars caused their relative similarity to change somewhat from that observed during the MDS study, the role of the MDS was simply to select parent stimuli that would not result in a degenerate morph space (e.g., if 2 parents were far more similar to one another than the other parents). Extensive pilot testing was conducted to determine the actual proportions of parents along each dimension when constructing the morph space itself (see Folstein et al. Forthcoming). Furthermore, because we counterbalanced which dimension was

relevant and which irrelevant across subjects, we are confident that the observed results reflect category learning and not any a priori similarity or dissimilarity between parents.

We created a 2D object space by morphing the parents in the following way: 2 parents (A and B) defined the X dimension and 2 parents (C and D) defined the Y dimension. Each continuous dimension was created by morphing between its 2 parents (i.e., a morphline between parents A and B defined the X dimension, and a morphline between parents C and D defined the Y dimension). A particular object in the 2D space (x,y) is created by morphing between image x along the X dimension and image y along the Y dimension. This technique for creating 2D morph spaces is adapted from Goldstone and colleagues (Goldstone and Steyvers 2001; Gureckis and Goldstone 2008). We sampled systematically from this space to obtain 64 objects (8×8) for use during category learning and 16 (4×4) additional objects for behavioral discrimination tests and for use in the main fMRI experiment. The precise position of these objects in the 2D space is shown in Supplementary Figure S1. The exact contribution of each parent to each morph object is available upon request. Morphs were produced using `gtkmorph` (`xmorph.sourceforge.net`).

Procedure

On the day of the fMRI scan, subjects first completed a category-learning task outside of the scanner, which was immediately followed by the fMRI scan, which in turn was immediately followed by a visual discrimination task outside of the scanner.

Category Learning

Immediately prior to scanning, subjects learned to categorize 2 novel brands of cars made by an imaginary manufacturer as either “Cags” or “Mons.” Critically, only 1D in the morph space was relevant for the learned boundary. Orientation of the category boundary (vertical vs. horizontal) was counterbalanced across subjects; if the category boundary was vertical, the horizontal dimension was relevant; if the category boundary was horizontal, the vertical dimension was relevant (for illustration, see Fig. 1). The categorization stimulus set included 64 cars sampled from the same stimulus space as the 16 stimuli used for the scanning task. The 8×8 stimulus set used for category learning is shown in Supplementary Figure S1 (filled circles). Subjects learned to categorize the stimuli over the course of 768 trials outside of the




Figure 1. Illustration of the 2D morphspace of cars used in the fMRI experiment. Two morphlines, each between a different set of parent cars, constitute the dimensions of the space (Parent A to Parent B morphline for the horizontal dimension, Parent C to Parent D morphline for the vertical dimension). The space was constructed by blending factorially between these two morphlines. Examples of relevant and irrelevant stimulus pairs are shown relative to a vertical category boundary. Note that the orientation of the category boundary was counterbalanced across subjects.

scanner, receiving corrective feedback after each trial. Each category learning trial proceeded as follows: On each trial, a fixation dot (500 ms) was followed by a car image (1500 ms). Subjects responded within 10s of the onset of the car image; if they did not respond within 10 s, they received a message asking them to respond a bit faster. Responses were followed by feedback (700 ms) consisting of the correct category or the car (“Cag” or “Mon”) and whether the response was correct or incorrect. Intertrial interval was 1000 ms.

Subjects categorized one more block of 64 cars inside the scanner, immediately prior to the scanning session.

Visual Discrimination

Immediately after scanning, subjects completed a same-different discrimination task to measure perceptual discriminability along the relevant versus irrelevant dimensions. There were a total of 12 discrimination blocks (384 total discrimination trials) per session. Each block consisted of all 24 possible different trials (i.e., all adjacent pairs) and 8 identical trials in a random order. All 16 possible identical trials were presented every 2 blocks. The “different” trials consisted of adjacent pairs, cars that differed by a single horizontal or vertical position in the space. Half of the different pairs differed along the relevant dimension, while the other half differed along the irrelevant dimension (Fig. 1, Supplementary Fig. S1). Each behavioral discrimination trial proceeded as follows: 700 ms fixation, 1500 ms sample car image, 300 ms black and white noise mask, 300 ms blank screen, the match car image, displaced slightly from center, which remained on screen until the participant responded or 5 s elapsed. Intertrial interval was 1 s.

Scanning Tasks

We used fMRI adaptation (Grill-Spector et al. 2006) to investigate the neural correlates of acquired distinctiveness during a task where pairs of cars had to be matched based simply on their location (Fig. 2). Under ideal conditions, this location task would make category membership entirely task irrelevant because it was not diagnostic of correct responses. We acknowledge, however, that attention to category-relevant dimensions might have carried over from categorization tasks performed prior to and during the scan. We address this issue in the discussion. If the relevant dimension of the visual representations is stretched as a consequence of category learning, then the BOLD response would be greater for pairs of objects that differed along the dimension previously relevant to the learned categories compared with pairs that differed along the irrelevant dimension (Fig. 1). Importantly, the choice of relevant dimension was counterbalanced across subjects, so that any observed effects of relevance cannot be attributed to physical differences between the stimulus pairs. Increased BOLD response for relevant pairs indicates greater release from neural adaptation and, in turn, greater neural discriminability for relevant pairs (Jiang et al. 2007). In separate runs, subjects also performed a match-to-category task on the same car pairs to assess the effect of active categorization. (Unfortunately, interpretation of the match-to-category task turned out to be complicated by an unanticipated confound, namely that the relevant pairs that cross the category boundary [“Relevant Cross” in Fig. 1] are associated with different responses, while all other pair types are associated with the same

response. In other words, subjects must respond “different” to pairs that cross the category boundary and “same” to all other pair types. We note here that relevant pairs that crossed the category boundary tended to elicit less activation than all other pair types. While this result is puzzling, it has been observed in other labs [Aguirre G, personal communication], and we do not attempt to interpret it here. The contrast relevant-same > irrelevant-same, which does not suffer from this confound, activated a lateral area of left extrastriate cortex [Talairach coordinates, TAL, Talairach and Tournoux 1988: -40, -81, -11] and a bilateral area of anterior cingulate cortex [TAL: 4, 31, 12; -4, 31, 11].)

Subjects were scanned 15–30 min after the category-learning task was completed. All subjects completed 4 runs of the match-to-location task and 4 runs of a match-to-category task, followed by localizer tasks for face, object, and spatial attention areas (the spatial attention localizer provided no result of interest and will not be discussed further). The location and category runs were ordered in groups of two, so that subjects performed 2 runs of one task followed by 2 runs of the second task until all 4 runs of both tasks were completed. The order of the 2 tasks was counterbalanced. The entire scanning session lasted about 1 h and 50 min.

All runs began with an 8 s fixation epoch. Each trial consisted of a pair of stimuli. The first stimulus appeared for 1.5 s followed by a noise mask for 0.5 s, followed by the second stimulus for 1 s, followed by a 3-s intertrial interval. The center of each stimulus was shifted very slightly above or below fixation. In the match-to-location task, subjects responded to the location of the second stimulus, indicating if it was identical or different from the first stimulus. In the match-to-category task, the participant indicated whether the second car was in the same category as the first.

Each match-to-location and match-to-category run included 64 stimulus pair presentations and an additional 8 null trials. Stimulus pairs were presented in random order with the constraint that both members of the pair had to be at least 2 “steps” (or “city blocks”) away in the 4×4 space from both members of the preceding pair. For any given stimulus pair, each member of the pair was presented first in half of the trials it appeared and above fixation in half the trials it appeared. All pairs were “same location” pairs on half of the trials they appeared. “Same location,” “different location,” “same category,” and “different category” trials each appeared with 50% probability. All positions in the stimulus space appeared with equal frequency during the scanning task. In this study, we considered only stimulus pairs whose members occupied adjacent positions within the stimulus space. Other nonadjacent pairs were included for purposes of balancing certain stimulus characteristics, but these pairs were not analyzed because they differed along both relevant and irrelevant dimensions, complicating the interpretation of any data they might provide. Further details about the stimuli presented during scanning are described in Supplementary Figure S2.

Object-sensitive regions of interest (ROIs) in the ventral stream were localized using a 1-back task presented over 2 runs in which subjects viewed 6 blocks per run of objects, faces, and scrambled objects. There were 16 stimuli, lasting 750 ms each, in each block, and blocks were presented in counterbalanced order. Subjects were instructed to press a button if they saw 2 images in a row that were identical. Each block contained 2 or 3 repeats. Each run lasted about 5 min.

Anatomical and Functional MRI

Whole-brain fMRI scans were performed on a 3-T Philips Intera Achieva MRI scanner using a gradient-echo echo-planar imaging pulse sequence (33 slices, 0.5 mm gap, $3 \times 3 \times 3$ mm voxel size; time repetition [TR] = 2 s, time echo [TE] = 35 ms). High resolution T_1 -weighted anatomical volumes were acquired prior to the functional runs using a 3D Turbo Field Echo acquisition (170 contiguous axial slices, $1 \times 1 \times 1$ mm voxel size, TR = 8 ms, TE = 3.7 ms).

Using Brain Voyager QX v1.10 (www.brainvoyager.com), data were subjected to 3D motion correction, temporal filtering (3 cycles/scan and high pass), and spatial smoothing (6-mm full-width at half-maximum Gaussian). In rare cases of motion exceeding 3 mm, trials 6–12 s prior to the motion were removed from the analysis and modeled in a separate confound predictor bin. All functional images were coregistered to the anatomical images and warped into standard Talairach space.




Figure 2. Illustration of the match-to-location task. Subjects judged whether pairs of identical or adjacent cars in the 4×4 grid were in the same location (slightly above or below fixation) or different locations. Stimuli spanned about 3.5×1.7 degrees of visual angle, both inside and outside of the scanner. Note that the displacement of the cars in the figure is not to scale, as cars were displaced by far less than the width of a car.

Statistical Analyses

Whole-brain contrasts were conducted using Brain Voyager's random effects-general linear model (GLM) procedure. For each subject, for each event in a given condition, the percent signal change for each voxel, starting with the second stimulus of the pair (the first stimulus was counterbalanced such that, across subjects, relevant and irrelevant pairs contained the same first stimuli), was correlated with a standard double-gamma hemodynamic function, which was time locked to the event. The beta weights resulting from this correlation were then used as the dependent measure in a second-level general linear model which included predictors for the following conditions: Relevant Cross, Relevant Same, Irrelevant Cross, Irrelevant Same, Identical, a single additional predictor for all other stimulus types, and predictors for each of 6 motion directions. For all whole-brain contrasts, an initial contrast with an uncorrected threshold of $P < 0.01$ was initially used to correct for multiple comparisons with Brain Voyager's Cluster Threshold Estimator (Forman et al. 1995; Goebel et al. 2006), which uses a bootstrapping procedure to calculate the minimum cluster size that yields a corrected threshold of $P < 0.05$. For ROI analyses, a beta weight was first calculated for each condition for each subject using the same canonical hemodynamic function as was used in the whole-brain analysis. The resulting beta weights were then used as the dependent measures in analyses of variance (ANOVAs).

While all analyses used a correlational approach, modeling BOLD time courses with canonical hemodynamic functions, all key figures include time courses calculated using event-related averaging. These reveal that time courses modeled using the GLM were unlikely to violate the assumptions of the statistical model. In all event-related averages, baseline is the average activation for the run and time zero is the onset of the second stimulus of the pair. The time course for the null condition was also subtracted from each condition.

Results

Behavior

Behavioral Performance outside the Scanner

Subjects categorization accuracy ranged from 80% to 93%. Category learning produced acquired distinctiveness: Following learning, replicating other studies (Goldstone 1994; Op de Beeck et al. 2003; Gureckis and Goldstone 2008) and our own prior behavioral findings with this stimulus space, discriminability was higher along the relevant dimension diagnostic of the learned categories than the irrelevant dimension (Fig. 3, 2×2 ANOVA with factors of Relevance and Boundary Crossing: Main effect of Relevance: $F_{1,19} = 10.9$, $P < 0.005$; all other effects not significant).

Behavioral Performance inside the Scanner

Average accuracy on a block of category learning performed in the scanner immediately prior to the functional scans was 87% (range 72–94%).

Sensitivity for detecting shifts in location in the match-to-location task, while not at ceiling, was well above chance in all stimulus conditions (mean d' 1.33, $t_{19} > 6.5$, $P_s < 0.0005$). Single factor ANOVAs conducted on reaction time or sensitivity, with 5 levels corresponding to each of the critical conditions (Relevant Cross, Irrelevant Cross, Relevant Same, Irrelevant Same and Identical; Table 1) were not significant (reaction time: $F_{4,76} = 0.637$, $P = 0.638$; d' : $F_{4,76} = 0.724$, $P = 0.58$). These analyses did not include 2 match-to-location runs for one subject that were removed from both behavioral and fMRI analyses due to near chance levels of performance.

fMRI Adaptation

Using data collected during the location task, we performed 2 contrasts to assess acquired distinctiveness specific to the




Figure 3. Results of the discrimination posttest administered after the fMRI scan. Sensitivity, measured using d' , was higher for relevant than irrelevant pairs and, as expected, there was no interaction between Relevance and Boundary Crossing. Error bars show 95% confidence intervals for the within-subject contrast. Note that d' in both conditions was significantly greater than a chance d' value of 0 ($t_{19} > 7$).

Table 1

Sensitivity (d') and reaction time for location detection task performed inside of the scanner. Ninety-five percentage of confidence intervals for main effect of task are shown in parentheses

	Relevant Cross	Relevant Same	Irrelevant Cross	Irrelevant Same	Identical
d' (95% CI)	1.226 (0.241)	1.300 (0.241)	1.358 (0.241)	1.289 (0.241)	1.482 (0.241)
Reaction time in seconds (95% CI)	1.018 (0.026)	1.005 (0.026)	0.989 (0.026)	1.005 (0.026)	1.007 (0.026)

relevant dimension. First, to isolate areas sensitive to an increase in discriminability along the entire relevant dimension, we performed a whole-brain contrast comparing all relevant pairs to all irrelevant pairs (Fig. 4a). This revealed activity in ventral visual cortex in the left anterior fusiform gyrus.

For additional power and to compare with ROIs analyzed by Jiang et al. (2007), we also performed the same contrast in functional ROIs in the ventral stream. To make strong claims about the sensitivity of particular functional areas (e.g., PPA, FFA, and LO) to a given manipulation, it is standard to define ROIs in individual subjects. However, our design requires a value for each subject because of the counterbalancing between the vertical and horizontal boundaries. For this reason, we could not use individually defined functional ROIs because they could not be defined in a subset of subjects. (To maintain counterbalancing, we would have had to remove up to 6 subjects from the analysis, which would have significantly decreased our statistical power.) We therefore used functionally defined ROIs based on the group-averaged localizer runs, which we will refer to neuroanatomically and by the functional contrast used to define them (Fig. 5). Beta weights for the conditions of Relevant Cross, Relevant Same, Irrelevant Cross, and Irrelevant Same were extracted from each of these ROIs and entered into a 2×2 ANOVA with factors Relevance and Boundary (whether the pair crossed or did not cross the middle of the space). This analysis revealed a global increase in neural discriminability along the relevant dimension in left lateral occipitotemporal cortex (objects > scrambled objects) and parahippocampal gyrus (objects > faces, Fig. 6). Main effects of Relevance, reflecting greater activation for relevant compared with irrelevant pairs were observed in both ROIs (left lateral occipital cortex:

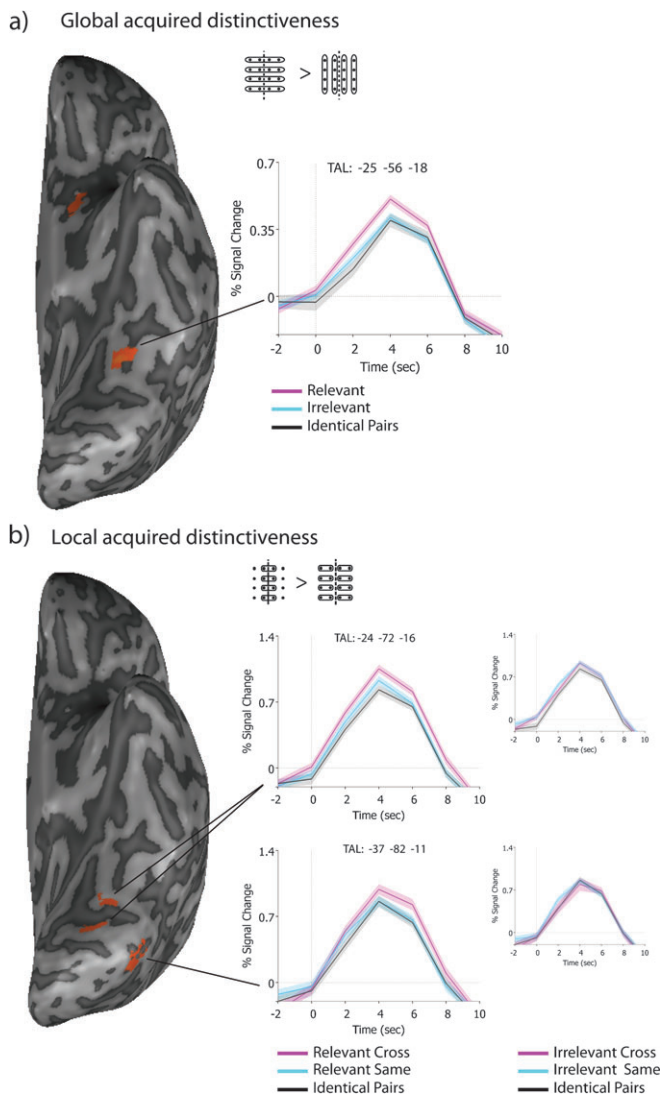


Figure 4. (a) Ventral view of the whole-brain comparison of all relevant pairs compared with all irrelevant pairs. Additional areas of activation are reported in Table 1. (b) Results for stimulus pairs that differ along the relevant dimension and are also in different categories ("Relevant Cross") compared with relevant same-category pairs ("Relevant Same")—additional activation was observed in retinotopic cortex and the cerebellum and an area of deactivation was observed in orbitofrontal cortex (Table 2). Large time courses are shown comparing Relevant Cross, Relevant Same, and Identical pairs, while small time courses illustrate that analogous pairs differing along the irrelevant dimension do not show the same effect. Color regions around the time-course lines represent standard errors.

$F_{1,19} = 4.85$, $MSe = 0.43$, $P = 0.04$; left parahippocampal gyrus: $F_{1,19} = 4.39$, $MSe = 0.13$, $P = 0.05$). Neither the main effect of Boundary nor the Relevance \times Boundary interaction were significant in either ROI (Boundary: left lateral occipital cortex: $F_{1,19} = 0.035$, $MSe = 0.639$, $P = 0.85$; left parahippocampal gyrus: $F_{1,19} = 0.92$, $MSe = 0.17$, $P = 0.35$; Relevance \times Boundary: left lateral occipital cortex: $F_{1,19} = 1.3$, $MSe = 0.96$, $P = 0.27$; left parahippocampal gyrus: $F_{1,19} = 0.36$, $MSe = 0.16$, $P = 0.56$). Equivalent analyses using deconvolution methods and directly analyzing event-related averages returned similar results (Supplementary Fig. S3). In contrast, we observed no effects of relevance in the right hemisphere ROIs (Supplementary Fig. S4).

The whole-brain contrast for relevant pairs greater than irrelevant pairs also revealed activation in regions outside of

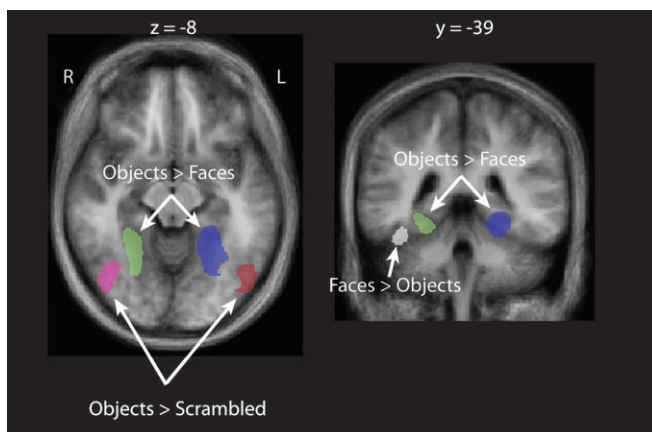


Figure 5. Object-selective ROIs were identified using standard functional localizer runs administered following the experimental runs. Each of the 2 localizer runs included blocks of objects, faces, and scrambled objects. Contrasts were performed at the group level and thresholds for each ROI were adjusted such that the ROI covered approximately 2000 mm³ in each hemisphere, with the exception of the parahippocampal gyrus (objects > faces), which covered approximately cubic 6000 mm. The contrast objects > scrambled objects revealed a cluster that included bilateral ventrolateral occipitotemporal cortex and the posterior fusiform sulcus (TAL: 42, -65, -9; -47, -67, -6). The contrast objects > faces revealed a large bilateral area along the parahippocampal gyrus (TAL: 25, -48, -10; -24, -48, -11). Finally, a face sensitive area in the right midfusiform was identified using the contrast faces > objects (TAL: 40, -42, -20).

visual cortex, including parts of the hippocampus, prefrontal cortex, and superior temporal gyrus (Table 2).

Our design also allowed us to perform other comparisons analogous to those in prior studies (Jiang et al. 2007; Gillebert et al. 2008). Again using a whole-brain contrast, we compared activity for relevant pairs that crossed the category boundary (Relevant Cross condition) with activity for relevant pairs that did not cross the boundary (Relevant Same condition). This revealed clusters of activation in left ventral and posterolateral occipital cortex in and/or near human area V4 in visual cortex (Wilms et al. 2010), just posterior to functionally localized Lateral Occipitotemporal Cortex (Fig. 4b, Table 3). (It is possible that the areas activated by this contrast could actually be sensitive to stimuli that crossed the middle of the stimulus space in either the relevant or irrelevant direction rather than specifically to stimuli that crossed the category boundary. To rule out this possibility, we contrasted all pairs that crossed the middle of the space with all pairs that did not cross the middle of the space [Supplementary Fig. S5 and Table S1]. This contrast did not activate any extrastriate areas.)

Discussion

We found that category learning causes stable increases in discriminability between neural populations in high-level visual areas that represent object dimensions that are relevant for categorization compared with dimensions that are irrelevant.

After subjects learned to categorize a 2D morphspace of cars, fMRI adaptation was measured during a match-to-location task that did not require categorization and where dimensions relevant to the learned categories were orthogonal to the location judgment. Stimulus pairs that differed along the relevant dimension adapted less than stimulus pairs that differed along the irrelevant dimension. Specifically, we observed reduced adaption along relevant object dimensions in a whole-brain

ROIs: location task




Figure 6. ROI results for the left hemisphere. Top row: beta weights, bottom row: ROI time courses. Beta weights for the conditions of Relevant Cross, Relevant Same, Irrelevant Cross, and Irrelevant Same were entered into a 2×2 ANOVA with factors Relevance and Boundary (crosses or does not cross the middle of the space). Main effects of Relevance, reflecting greater activation for relevant compared with irrelevant pairs were observed in 2 ROIs within the left hemisphere: lateral occipitotemporal cortex (defined by the contrast objects > scrambled objects), parahippocampal gyrus (defined by the contrast objects > faces) ($F_{S1,19} > 4$, $P_s < 0.05$). Neither the main effect of Boundary nor the Relevance \times Boundary interaction was significant in any of these ROIs. These results are consistent with global acquired distinctiveness for relevant dimensions.

analysis in a region near the midfusiform gyrus and in functionally localized object-selective ROIs. Our observed difference in neural discriminability following category learning mirrors the difference in perceptual discriminability we observed behaviorally (see also Folstein et al. Forthcoming). Past research that failed to find effects of category learning on representations of relevant dimensions in the visual system failed to report any effects of category learning on perceptual discrimination, which seems a logical prerequisite (e.g., Jiang et al. 2007).

Relatively few studies have compared neural modulation of relevant versus irrelevant dimensions of multidimensional object spaces following category learning. To date, the only studies to find selective neural sensitivity to a relevant dimension have done so based on data collected while subjects were actively categorizing objects (Sigala and Logothetis 2002; Li et al. 2007; De Baene et al. 2008). Other results, using paradigms that do not map exactly onto a comparison of relevant versus irrelevant dimensions, also suggest that visual cortex represents some category information during active categorization (Meyers et al. 2008). Such observed effects might have been due to flexible top-down dimensional weighting of object representations applied only during active categorization (Nosofsky 1984; Kruschke 1992) rather than to stable

changes in object representations (Gauthier and Palmeri 2002). Perhaps outside of active categorization, no such dimensional stretching applies, an interpretation consistent with studies reporting no category-specific increases in neural sensitivity in object-sensitive areas during noncategorization tasks (Jiang et al. 2007; Gillebert et al. 2008). By contrast, our results reveal that even with complex objects of the sort used in those studies, when behavioral evidence for acquired distinctiveness is obtained, an analogous effect in the visual system is obtained, even during a location judgment. Our evidence is compatible with theories that suggest that category learning can cause representations in visual cortex to become stably sensitized to relevant object dimensions but incompatible with theories where dimensional relevance has no systematic influence on object representations (Riesenhuber and Poggio 1999; Serre et al. 2007). Because some common kinds of methodological procedures used to create morphspaces can limit modulation of category-relevant dimensions (Folstein et al. Forthcoming), the field has likely underestimated the impact of category learning on neural visual representations.

Our discussion so far has emphasized the stability of the changes to visual cortex as a consequence of category learning, but much work remains to be done to determine whether

Table 2

Match-to-location task: relevant > irrelevant

	TAL		
	X	Y	Z
Relevant > Irrelevant			
Frontal areas			
R. precentral gyrus	3.7	-22	69
R. middle/superior frontal gyrus	29	34	41
R. piriform cortex	22	7	-10
L. precentral gyrus	-19	-39	61
Parietal areas			
R. postcentral gyrus	63	-20	20
R. postcentral gyrus	20	-36	60
Temporal areas			
R. middle temporal gyrus	46	-65	16
R. superior temporal gyrus	41	-22	-2
R. posterior parahippocampal gyrus	8.5	-38	0.94
L. fusiform gyrus	-29	-56	-17
Insular areas			
R. posterior insula	36	-11	-2.3
R. posterior insula	36	-30	16
Basal ganglia/limbic areas			
L. hippocampus	-20	-22	-11
L. amygdala/basal forebrain	-14	-6	-11
R. nucleus accumbens	9	4	-4
Other areas			
L. cerebellum	-22	-57	-20

Note: R, right; L, left.

Table 3

Match-to-location task: Relevant Cross > Relevant Same and Relevant Same > Relevant Cross

	TAL		
	X	Y	Z
Relevant Cross > Relevant Same			
Occipital areas			
L. extrastriate cortex (V4)	-25	-73	-18
L. lateral extrastriate cortex (V4)	-37	-82	-12
L. striate cortex	-1	-91	-10
R. striate cortex	4	-91	-10
Other areas			
L. cerebellum	-13	-66	-23
	-10	-72	-35
	-6	-72	-24
R. cerebellum	26	-62	-25
Relevant Same > Relevant Cross			
Frontal areas			
R. orbital inferior frontal gyrus	30	23	-14

these effects could be driven, even in part, by top-down modulation, even outside of active categorization. In our study, pairs differing along the relevant dimension also caused greater activation than irrelevant pairs in several areas outside the ventral stream, including dorsolateral prefrontal cortex (DLPFC), middle and superior temporal gyrus, hippocampus, and reward-related areas such as the amygdala and nucleus accumbens. One potential explanation for activations in these nonvisual areas might be that changes between the first and second stimulus along the relevant dimension “caught the subjects’ attention,” causing a kind of orienting response that activates visual cortex. However, one finding in our data argues against this interpretation. If subjects oriented to shape changes along the relevant dimension, one might expect this to take attention away from the location task, slowing reaction time or decreasing accuracy. In fact, there were no differences between the critical conditions in either reaction time or accuracy.

Another possible explanation for our findings, distinct from the orienting explanation, is that subjects were covertly

categorizing the stimuli while also doing the location task. While this possibility is not easy to rule out, we note that the location task was by no means easy. Performance was well below ceiling ($d' 1.33$); while above chance, this suggests that subjects had to attend to the primary location task. Given the large cost associated with performing 2 even relatively simple tasks concurrently (Pashler 1994), it seems unlikely that subjects were judging location and categorizing, especially given that both were difficult and that categorization was completely uninformative regarding location.

A related, but more subtle alternative explanation is that attention to relevant dimensions carried over to the location task from the temporally proximal categorization runs even though the location task used a dimension irrelevant to the learned categories. While this top-down explanation is difficult to rule out, a persistent attentional bias for category-relevant dimensions during an orthogonal task could be seen as one possible stable form of acquired distinctiveness. One open question is how long this stable attentional bias can persist. Ongoing work is assessing how long both the neural and behavioral signatures of acquired distinctiveness can persist and whether their locus is bottom-up or top-down.

Finally, it is possible that the nonvisual areas we observed could have been engaged by visual cortex in a bottom-up fashion and potentially under the modulation of the effects detected in the visual system. Several of these areas could represent category-relevant dimensions such as the degree to which a stimulus is associated with a reward or particular outcome. For instance, category-sensitive cells have been observed in the medial temporal lobe (Kreiman et al. 2000; Hampson et al. 2004) and the superior temporal gyrus (van der Linden et al. 2010). Finally, even if the activations we observed in the visual system were the result of top-down modulation, our results would nonetheless challenge the interpretation of past studies that have not observed acquired distinctiveness effects in the visual system whether subjects were actively categorizing or not (Jiang et al. 2007). One next step in evaluating the possibility that these effects are bottom up might involve investigating how long these behavioral and fMRI effects endure without any intervening category learning.

In addition, this work significantly extends that of Sigala and Logothetis (2002) and De Baene et al. (2008), both of whom used stimuli with relatively simple dimensions (spatially separated parts and curvature/aspect ratio, respectively). Our study demonstrates that learning categories defined by much more complex dimensions can not only produce acquired distinctiveness (Goldstone and Steyvers 2001; Gureckis and Goldstone 2008) but can also elicit selective enhancement of relevant dimensions in visual cortex. Indeed, previous work, including recent behavioral work in our own lab, suggests that the effect of category learning goes beyond acquired distinctiveness of preexisting dimensions, to the possible creation of relevant features and dimensions useful to novel category learning (Schyns et al. 1998; Goldstone and Steyvers 2001; Folstein et al. Forthcoming). It is possible that the effects we observed in visual cortex are a neural correlate of feature creation induced by category learning.

We should emphasize that several of the studies reviewed in this paper that reported effects of categorization in the visual system did not compare relevant and nonrelevant dimensions to assess acquired distinctiveness. Instead, many compared neural sensitivity to pairs of objects that crossed a category

boundary to those that did not (see also Li et al. 2009), an effect more local to the boundary region, which is sometimes called categorical perception (Harnad 1987). Our results are the first to suggest that different parts of the visual system show a more global form of acquired distinctiveness along an entire relevant dimension compared with a more local form of categorical perception across a category boundary; local effects, likely in extrastriate area V4, were more posterior to the fusiform regions showing global effects. Unfortunately, our neural finding is qualified by the fact that we did not observe a behavioral effect across the category boundary—future work with longer category training can investigate the possibility that neural effects across the boundary are associated with behavioral signatures of categorical perception.

Several studies have emphasized the role of prefrontal cortex in representing categories and have suggested that the ventral stream primarily represents shape, irrespective of learned categories. In some cases, PFC cells represent category in a flexible task-specific manner (Cromer et al. 2010), while in other cases, PFC neurons appear to show some degree of task-independent coding of categories, representing one category boundary while the animal is categorizing according to a different boundary (Roy et al. 2010). Our study extends these findings by revealing task-irrelevant modulation of human prefrontal cortex by relevant object dimensions. But prefrontal cortex does not appear to be the sole locus of category learning and recent work suggests that it may not be necessary: for instance, animals can learn categories even when PFC is lesioned (Minamimoto et al. 2010).

While the role of the PFC in categorization has been prominent in studies where categorization has little or no effect in the visual system (Seger and Miller 2010), it is interesting to consider its role in a case where category learning does tune visual areas to relevant object dimensions. The active area of DLPFC we observed was similar (Euclidean distance in Talairach space = 12 mm) to an area previously observed to encode category information during category learning (Li et al. 2007). In addition, right DLPFC has been engaged by various contrasts in the context of active categorization (Seger et al. 2000; Vogels et al. 2002; Koenig et al. 2005; Cincotta and Seger 2007). All of these findings are consistent with some sensitivity for DLPFC to relevant dimensions. But, interestingly, active DLPFC areas in these studies tend to be dorsal to the prefrontal area observed by Jiang et al. (2007) that responded more to boundary crossing pairs during a match-to-category task. Unlike our study, Jiang et al. (2007) trained subjects from the beginning on a match-to-category task rather than on a category-to-response mapping task, which is most typical of other fMRI categorization studies. It is interesting to speculate that dorsal regions of DLPFC represent relevant object features in the service of linking them with category responses while more ventral regions of DLPFC represent them in the service of holding them in working memory. This is consistent with a recent proposal by O'Reilly (2010) for a “how” to “what” dorsal to ventral continuum in PFC (see also Miller and Cohen 2001).

In closing, we demonstrate improvements both in perceptual discriminability and in discriminability of the neural representations of relevant dimensions when categories are irrelevant. While explicit categorization decisions may critically involve nonvisual areas, our results demonstrate that when category learning improves our ability to perceive objects, this reflects changes in visual cortical representations.

Supplementary Material

Supplementary material can be found at: <http://www.cercor.oxfordjournals.org/>

Funding

National Institutes of Health (grant 1 F32 EY019445-01, grant 2 RO1 EY013441-06A2, P30-EY008126); Temporal Dynamics of Learning Center; National Science Foundation (grant SBE-0542013).

Notes

We thank Jeffrey Schall, Randolph Blake, René Marois, and members of the Gauthier and Palmeri laboratories for comments on this manuscript. We also thank Jascha Swisher and Chris Asplund for suggestions on data analysis. *Conflict of Interest:* None declared.

References

- Cincotta CM, Seger CA. 2007. Dissociation between striatal regions while learning to categorize via feedback and via observation. *J Cogn Neurosci.* 19:249–265.
- Corbetta M, Miezin FM, Dobmeyer S, Shulman GL, Petersen SE. 1991. Selective and divided attention during visual discriminations of shape, color, and speed: functional anatomy by positron emission tomography. *J Neurosci.* 11:2383–2402.
- Cromer JA, Roy JE, Miller EK. 2010. Representation of multiple, independent categories in the primate prefrontal cortex. *Neuron.* 66:796–807.
- De Baene W, Ons B, Wagemans J, Vogels R. 2008. Effects of category learning on the stimulus selectivity of macaque inferior temporal neurons. *Learn Mem.* 15:717–727.
- Drucker DM, Kerr WT, Aguirre GK. 2009. Distinguishing conjoint and independent neural tuning for stimulus features with fMRI adaptation. *J Neurophysiol.* 101:3310–3324.
- Folstein JR, Gauthier I, Palmeri TJ. Forthcoming. How category learning affects object discrimination: not all morphspaces stretch alike. *J Exp Psychol Learn Mem Cogn.* doi: 10.1037/a0025836.
- Forman SD, Cohen JD, Fitzgerald M, Eddy WF, Mintun MA, Noll DC. 1995. Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold. *Magn Reson Med.* 33:636–647.
- Freedman DJ, Miller EK. 2008. Neural mechanisms of visual categorization: insights from neurophysiology. *Neurosci Biobehav Rev.* 32:311–329.
- Freedman DJ, Riesenhuber M, Poggio T, Miller EK. 2003. A comparison of primate prefrontal and inferior temporal cortices during visual categorization. *J Neurosci.* 23:5235–5246.
- Gauthier I, Palmeri TJ. 2002. Visual neurons: categorization-based selectivity. *Curr Biol.* 12:R282–R284.
- Gillebert CR, Op de Beeck HP, Panis S, Wagemans J. 2008. Subordinate categorization enhances the neural selectivity in human object-selective cortex for fine shape differences. *J Cogn Neurosci.* 21:1054–1064.
- Goebel R, Esposito F, Formisano E. 2006. Analysis of functional image analysis contest (FIAC) data with brainvoyager QX: from single-subject to cortically aligned group general linear model analysis and self-organizing group independent component analysis. *Hum Brain Mapp.* 27:392–401.
- Goldstone RL. 1994. Influences of categorization on perceptual discrimination. *J Exp Psychol Gen.* 123:178–200.
- Goldstone RL, Steyvers M. 2001. The sensitization and differentiation of dimensions during category learning. *J Exp Psychol Gen.* 130:116–139.
- Grill-Spector K. 2003. The neural basis of object perception. *Curr Opin Neurobiol.* 13:159–166.
- Grill-Spector K, Henson R, Martin A. 2006. Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn Sci.* 10:14–23.
- Gureckis TM, Goldstone RL. 2008. The effect of the internal structure of categories on perception. In: Love BC, McRae K, Sloutsky VM,

- editors. Proceedings of the 30th Annual Conference of the Cognitive Science Society; 2008 Jul 23–26. Austin (TX): Cognitive Science Society. p 1876–1881.
- Gureckis TM, James TW, Nosofsky RM. 2011. Re-evaluating dissociations between implicit and explicit category learning: an event-related fMRI study. *J Cogn Neurosci*. 23:1697–1709.
- Hampson RE, Pons TP, Stanford TR, Deadwyler SA. 2004. Categorization in the monkey hippocampus: a possible mechanism for encoding information into memory. *Proc Natl Acad Sci U S A*. 101:3184–3189.
- Harnad SR. 1987. *Categorical perception: the groundwork of cognition*. Cambridge: Cambridge University Press, p. 599.
- Hockema SA, Blair MR, Goldstone RL. 2005. Differentiation for Novel Dimensions. In: Bara B, Barsalou L, Bucciarelli M, editors. Proceedings of the Twenty-seventh Annual Conference of the Cognitive Science Society; 2005 Jul 21–23. Hillsdale (NJ): Lawrence Erlbaum Associates. p. 953–958.
- Jiang X, Bradley E, Rini RA, Zeffiro T, Vanmeter J, Riesenhuber M. 2007. Categorization training results in shape- and category-selective human neural plasticity. *Neuron*. 53:891–903.
- Koenig P, Smith EE, Glosser G, DeVita C, Moore P, McMillan C, Gee J, Grossman M. 2005. The neural basis for novel semantic categorization. *Neuroimage*. 24:369–383.
- Kreiman G, Koch C, Fried I. 2000. Category-specific visual responses of single neurons in the human medial temporal lobe. *Nat Neurosci*. 3:946–953.
- Kruschke JK. 1992. ALCOVE: an exemplar-based connectionist model of category learning. *Psychol Rev*. 99:22–44.
- Li S, Mayhew SD, Kourtzi Z. 2009. Learning shapes the representation of behavioral choice in the human brain. *Neuron*. 62:441–452.
- Li S, Ostwald D, Giese M, Kourtzi Z. 2007. Flexible coding for categorical decisions in the human brain. *J Neurosci*. 27:12321–12330.
- Meyers EM, Freedman DJ, Kreiman G, Miller EK, Poggio T. 2008. Dynamic population coding of category information in inferior temporal and prefrontal cortex. *J Neurophysiol*. 100:1407–1419.
- Miller EK, Cohen JD. 2001. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci*. 24:167–202.
- Minamimoto T, Saunders RC, Richmond BJ. 2010. Monkeys quickly learn and generalize visual categories without lateral prefrontal cortex. *Neuron*. 66:501–507.
- Nosofsky RM. 1984. Choice, similarity, and the context theory of classification. *J Exp Psychol Learn Mem Cogn*. 10:104–114.
- Nosofsky RM. 1986. Attention, similarity, and the identification-categorization relationship. *J Exp Psychol Gen*. 115:39–61.
- Nosofsky RM, Little DR, James TW. 2012. Activation in the neural network responsible for categorization and recognition reflects parameter changes. *Proc Natl Acad Sci U S A*. 109:333–338.
- Notman LA, Sowden PT, Özgen E. 2005. The nature of learned categorical perception effects: a psychophysical approach. *Cognition*. 95:B1–B14.
- O'Reilly RC. 2010. The what and how of prefrontal cortical organization. *Trends Neurosci*. 33:355–361.
- Op de Beeck HP, Wagemans J, Vogels R. 2003. The effect of category learning on the representation of shape: dimensions can be biased but not differentiated. *J Exp Psychol Gen*. 132:491–511.
- Panis S, Vangeneugden J, Op de Beeck HP, Wagemans J. 2008. The representation of subordinate shape similarity in human occipito-temporal cortex. *J Vis*. 8:9.1–9.15.
- Pashler H. 1994. Dual-task interference in simple tasks: data and theory. *Psychol Bull*. 116:220–244.
- Reber PJ, Gitelman DR, Parrish TB, Mesulam MM. 2003. Dissociating explicit and implicit category knowledge with fMRI. *J Cogn Neurosci*. 15:574–583.
- Riesenhuber M, Poggio T. 1999. Hierarchical models of object recognition in cortex. *Nat Neurosci*. 2:1019–1025.
- Roy JE, Riesenhuber M, Poggio T, Miller EK. 2010. Prefrontal cortex activity during flexible categorization. *J Neurosci*. 30:8519–8528.
- Sawamura H, Orban GA, Vogels R. 2006. Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the fMRI adaptation paradigm. *Neuron*. 49:307–318.
- Schyns PG, Goldstone RL, Thibaut JP. 1998. The development of features in object concepts. *Behav Brain Sci*. 21:1–17; discussion 17–54.
- Seeger CA, Miller EK. 2010. Category learning in the brain. *Annu Rev Neurosci*. 33:203–219.
- Seeger CA, Poldrack RA, Prabhakaran V, Zhao M, Glover GH, Gabrieli JD. 2000. Hemispheric asymmetries and individual differences in visual concept learning as measured by functional MRI. *Neuropsychologia*. 38:1316–1324.
- Serre T, Oliva A, Poggio T. 2007. A feedforward architecture accounts for rapid categorization. *Proc Natl Acad Sci U S A*. 104:6424–6429.
- Sigala N, Logothetis NK. 2002. Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature*. 415:318–320.
- Talairach J, Tournoux P (1988). *Co-planar stereotaxic atlas of the human brain*. New York: Thieme.
- van der Linden M, van Ruennout M, Idefrey P. 2010. Formation of category representations in superior temporal sulcus. *J Cogn Neurosci*. 22:1270–1282.
- Vogels R, Sary G, Dupont P, Orban GA. 2002. Human brain regions involved in visual categorization. *Neuroimage*. 16:401–414.
- Wilms M, Eickhoff SB, Homke L, Rottschy C, Kujovic M, Amunts K, Fink GR. 2010. Comparison of functional and cytoarchitectonic maps of human visual areas V1, V2, V3d, V3v, and V4(v). *Neuroimage*. 49:1171–1179.