

Dynamic Norm-based Encoding for Unfamiliar Shapes in Human Visual Cortex

Sven Panis, Johan Wagemans, and Hans P. Op de Beeck

Abstract

■ Previous studies have argued that faces and other objects are encoded in terms of their deviation from a class prototype or norm. This prototype is associated with a smaller neural population response compared with nonprototype objects. However, it is still unclear (1) whether a norm-based representation can emerge for unfamiliar or novel object classes through visual experience at the time scale of an experiment and (2) whether the results from previous studies are caused by the prototypicality of a stimulus, by the physical properties of individual stimuli independent from the stimulus distribution, and/or by the trial-to-trial adaptation. Here we show with a combined behavioral and event-related fMRI study in humans that a short amount of visual

experience with exemplars from novel object classes determines which stimulus is represented as the norm. Prototypicality effects were observed at the behavioral level by behavioral asymmetries during a stimulus comparison task. The fMRI data revealed that class exemplars closest to the prototypes—the perceived average of each class—were associated with a smaller response in the anterior part of the visual object-selective cortex compared with other class exemplars. By dissociating between the physical characteristics and the prototypicality status of the stimuli and by controlling for trial-to-trial adaptation, we can firmly conclude for the first time that high-level visual areas represent the identity of exemplars using a dynamic, norm-based encoding principle. ■

INTRODUCTION

Primates are able to recognize objects at various levels, often referred to as superordinate-level classification (e.g., tools vs. animals), basic-level categorization (e.g., monkey vs. dog), and subordinate-level identification (e.g., “is this my dog?”) (Logothetis & Sheinberg, 1996; Nosofsky, 1986; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). Subordinate-level identification is very well developed in primates, with face recognition being a primary example, but the underlying neural representations are not completely understood. Here we studied the role of visual experience for the emergence of a specific coding scheme for representing the differences among exemplars of visually homogeneous (basic-level) categories.

Recently, it has been proposed that adaptive coding mechanisms, which are widely used to code simple low-level sensory properties, also underlie the high-level coding of face identity (Rhodes & Jeffery, 2006). Specifically, they propose that “facial identity is coded by pairs of neural populations that are adaptively tuned to above-average and below-average values, respectively, of each dimension in face-space (see Figure 1). The values of each face on each dimension of face-space are signaled by the relative activation of the paired populations, with equal activation signaling average values” (p. 2984).

An adaptive coding mechanism (1) is based on neurons that dynamically adapt to the prevailing (average or prototypical) stimulus values or statistics in the environment, (2) is neurally and computationally efficient, and (3) alerts the system by focussing resources on uncommon or non-average inputs (Barlow, 1990). Thus, discriminating between visually similar faces is possible if the average face, which lies at the center of a neural face space, functions as a norm or prototype against which individuating (identity) information is coded (Rhodes & Jeffery, 2006; Leopold, O’Toole, Vetter, & Blanz, 2001).

Evidence for such norm-based representations has been provided using behavioral paradigms measuring similarity ratings, aftereffects, perceptual classifications, same-different decisions, and so forth (Anderson & Wilson, 2005; Leopold, Rhodes, Muller, & Jeffery, 2005; Op de Beeck, Wagemans, & Vogels, 2003a; Leopold et al., 2001; Lee, Byatt, & Rhodes, 2000; Suzuki & Cavanagh, 1998; Nosofsky, 1991; Rhodes, Brennan, & Carey, 1987), human fMRI studies (Loffler, Yourganov, Wilkinson, & Wilson, 2005), and extracellular recordings in monkeys (De Baene, Premereur, & Vogels, 2007; Leopold, Bondar, & Giese, 2006; Kayaert, Biederman, Op de Beeck, & Vogels, 2005). Individual neurons represent how, and how much, an exemplar of a class deviates from the class norm or prototype. The population response, that is, the response averaged across all involved neurons, also represents the deviation from the class prototype, which is associated with the lowest population response. Importantly, several of these studies used other

University of Leuven (K.U. Leuven), Belgium




Figure 1. Adaptive coding model (adapted from Rhodes & Jeffery, 2006; see also Tsao & Freiwald, 2006). For each dimension x used to discriminate exemplars, there are two populations of neurons. Pool 1 (dotted curve) codes below-average values, and Pool 2 (black curve) codes above-average values. Average values are coded implicitly by equal activation of the two populations (vertical line). Exposure to an exemplar with a high value on dimension x will adapt Pool 2 neurons (vertical arrows and dashed curve) and shift the perceived average (dashed vertical line) toward the adapting exemplar (horizontal arrow).

shape stimuli than faces (De Baene et al., 2007; Kayaert, Biederman, Op de Beeck, et al., 2005; Op de Beeck et al., 2003a), suggesting that norm-based encoding might be a general principle by which different exemplars from a basic level category are represented.

However, none of these reports has explicitly investigated how such a norm-based encoding scheme is induced and whether there is an effect of visual experience. Although adaptive tuning to stimulus statistics, hence a strong effect of visual experience is an inherent feature of the adaptive coding model as described earlier, the aforementioned experiments have typically studied norm-based coding as a given and have taken for granted that the “norm” is a fixed stimulus.

A simple observation tells us that this cannot be true. Tsao and Freiwald (2006) compared the prototypes used in different studies on face space (Rhodes & Jeffery, 2006; Anderson & Wilson, 2005; Leopold et al., 2001) and concluded that these prototypes have a very different appearance: “The ‘average face’ used in three studies of the face adaptation illusion look quite different from each other..., yet all three studies argue for a special role for their average face” (p. 393). This suggests that the stimulus that is seen as the norm or prototype is not fixed, even not for very well known objects such as faces for which the prototype could have been firmly established throughout development or even be innate. It should thus be possible to reverse the role of stimuli as being a prototype or not by changing the distribution of stimuli, maybe already during

the course of an experiment. This is what we have done in the present study.

At the neural level, we used the previously described differences in neural population response (Loffler et al., 2005) as an index for whether a stimulus is processed as a prototype or not. At the behavioral level, we analyzed the asymmetry of similarity judgments (Tversky, 1977). Several previous studies have proposed that asymmetries in a successive stimulus comparison task can be used as an index that some stimuli have a special status in the representation of stimulus differences (Op de Beeck et al., 2003a; Nosofsky, 1991). The underlying notion is the concept of stimulus bias, that is, a bias to perceive or remember certain stimuli. For example, “natural prototypes” defined by Rosch (1973) as elements around which many perceptual categories found in the natural world are organized have been described as a clear-cut example of stimulus bias (see Nosofsky, 1991, p. 114–115), and asymmetric similarity judgments related to the prototypicality of shapes have been described before (Op de Beeck et al., 2003a). The direction of the empirically observed and theoretically predicted (Nosofsky, 1991) asymmetry is as follows: If in a trial the special (or biased) stimulus comes first, then it will be perceived as more different from the other stimulus in the trial than when the special stimulus comes second. If the two stimuli in a trial have the same status (e.g., no prototypes), then no asymmetry is expected.

Once we consider the possibility that these prototype effects are induced by visual experience, we have to consider at which time scale these effects play. A longer term, adaptive effect of norm-based coding that builds up across many trials would not decrease the value of the concept of norm-based coding and the way it is measured in neurophysiological studies. However, some previous, as yet unpublished, reports have criticized, even dismissed, the neurophysiological results on norm-based coding for faces (Leopold et al., 2006; Loffler et al., 2005) because they did not take into account the possible effects of very short-term trial-to-trial adaptation (N. Davidenko, D. Remus, and K. Grill-Spector, *Society for Neuroscience Abstracts*, 19.5, 2008). The criticism is that very short-term effects, such as adaptation from the immediately preceding trial, were not ruled out. Since at this moment the only published fMRI study on norm-based encoding (Loffler et al., 2005) used a block design, this alternative explanation in terms of short-term adaptation cannot be ruled out yet.

Here, we wanted to establish (i) that a norm-based coding can be so strongly modulated by visual experience during the course of an experiment so that the physical stimulus properties or identity can be dissociated from the stimulus’ status as a prototype, even more, that the status as prototype or extreme class example can be reversed between groups of participants simply by the type of visual experience, and (ii) that these effects are not due to short-term trial-to-trial adaptation and build up gradually across intervening trials.

Here, we present the first study that dissociates between the prototypicality status of an exemplar and the other pre-existing biases associated with the shape of this exemplar. Stimuli were well-controlled geometric shapes that were located at equal physical and perceived distances in a (one-dimensional) shape space. The distribution of exemplars was changed across subjects so that the class prototype in one subject was a nonprototype in another subject. In other words, the prototypicality status of an exemplar could only be inferred by taken into consideration its similarity to the particular exemplars that were shown.

Furthermore, we designed the fMRI experiment in such a way to avoid adaptation between successive trials as much as possible. In contrast to all previous fMRI studies on this topic, we used an event-related design with one image per trial, and we had a relatively long time interval between successive images. We also verified empirically that no adaptation occurred specifically between successive images.

Biasing effects of prototypicality were evident at the behavioral level from behavioral asymmetries during a temporal stimulus comparison task with direct similarity ratings. Furthermore, the class exemplar closest to the perceived average was associated with a smaller fMRI BOLD response compared with nonprototypical class exemplars in object-selective visual cortex, more specifically in the fusiform gyrus. This finding supports the hypothesis that visual exposure to the distribution of novel object exemplars can strongly modulate and even reverse between groups of participants, which stimulus is represented as the class prototype in a higher level region of the cortical object vision pathway.

METHODS

Subjects

Thirty subjects (15 men, aged 20–34 years) participated in this experiment. Written informed consent was obtained from each subject, and they were paid €60 for their par-

ticipation (3 hours in total). All subjects participated in the two phases of the experiment (behavior and fMRI).

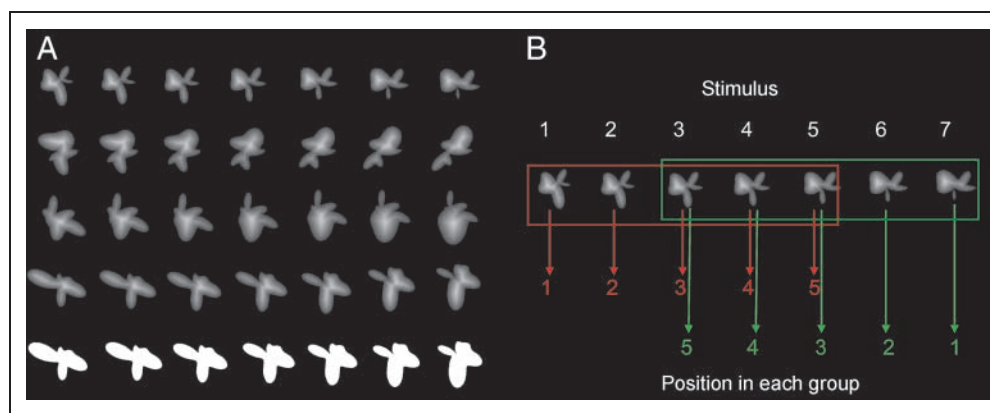
Ethics Statement

All procedures were approved by the relevant ethical boards, that is, the ethical committee of the Faculty of Psychology and Educational Sciences (K.U. Leuven) and the committee for medical ethics of the University Hospital. Written informed consent was obtained from each participant.

Stimuli

Stimuli were constructed using a parameterization in terms of radial frequency components as described before (Op de Beeck, Wagemans, & Vogels, 2001, 2003b). The total stimulus set included four classes of seven object exemplars. The seven exemplars from each class were located on a (one-dimensional) line in an originally two-dimensional shape space with dimensions representing the amplitude of two radial frequency components. The stimuli were gray surfaces presented on a black background (256×256 pixels). For 14 participants, stimuli were presented as uniform gray surfaces/silhouettes on a black background. For the other 16 participants, shapes were rendered with a grassfire texture on the basis of the shape skeleton (Blum, 1973), giving the stimuli the appearance of three-dimensional objects (see Figure 2A). Previous studies of object-selective regions in humans and monkeys have found selective responses with shaded three-dimensional objects (Op de Beeck, Torfs, & Wagemans, 2008) as well as with two-dimensional silhouettes or contours (Drucker & Aguirre, 2009; Kayaert, Biederman, & Vogels, 2005; Op de Beeck et al., 2001). Nevertheless, Georgieva, Todd, Peeters, and Orban (2008) found differences in the overall response strength in human lateral occipital complex (LOC), with a much stronger response to shaded objects, which raises the possibility that the two types of stimuli are represented

Figure 2. Stimuli and experimental design logic. (A) Four stimulus classes with all seven exemplars in each class rendered with the grassfire texture. The last row shows the silhouette versions of one class. (B) Each participant was presented with either the first five (red box) or the last five (green box) exemplars from a class. The assignment of stimulus exemplars to stimulus conditions (the positions shown at the bottom) was done in such a way that stimulus exemplar 3 was the prototype (Position 3) for the red group, but an extreme exemplar (Position 5) for the green group. Likewise, Stimulus 5 was the prototype (Position 3) for the green group but an extreme exemplar (Position 5) for the red group.



differently. We tested this by including this difference as a between-subject factor, with one subject group tested with silhouettes and the other group with shaded stimuli. This between-subject manipulation had no effect on the response strength in LOC, nor did it interact with the effects of interest (see Results).

We checked that the one-dimensional parametric manipulation also results in a one-dimensional physical change in the actual images. We computed the pixel-wise differences between all images in each object class (Op de Beeck et al., 2001; Grill-Spector et al., 1999). For each pair of images, we computed the difference in each pixel (maximal difference = 1), squared it, summed it across all pixels, took the square root of this sum, and normalized the resulting number by the square root of the number of pixels. The resulting difference matrix is shown in Table 2A and illustrates several points. First, the differences between successive images are very constant, always at 0.32. Second, and more in general, physical pixel-based differences conform very closely to what one would expect on the basis of a one-dimensional parametric variation with constant steps (e.g., stimuli placed at coordinates 0, 1, 2, 3, and 4, on the one dimension). In fact, the correlation between these parametric differences and pixel-based differences was very high, $r = .94$. Thus, the parametrically controlled shapes result in a good control of physical pixel-based variability.

Stimulus presentation and response registration were controlled by a laptop running the PsychToolbox (Brainard, 1997; Pelli, 1997) in Matlab (MathWorks, Natick, MA). Stimuli were shown on a CRT monitor during training (resolution = 1024×768 pixels, refresh rate = 60 Hz) and projected into a mirror in front of the head during scanning by means of a liquid crystal display projector (1280×1024 pixels; Barco 6300; Barco, Kortrijk, Belgium). Stimuli, maximum 8 visual degrees in width/height, were presented around the center of the screen with a random position offset from the center of the screen of maximally 2 visual degrees.

Experimental Design

To manipulate prototypicality while controlling for physical differences, we designed the experiment as follows (see Figure 2B). The stimulus set included four object classes with exemplars varying on one complex shape dimension (Op de Beeck et al., 2001, 2003b) (Figure 2A). There were seven stimuli in each class, but only five were shown to an individual subject. So, we had five stimulus conditions per subject, and the assignment of stimuli to stimulus conditions was changed across subjects. Stimulus Condition 3 is the prototype and stimulus Conditions 1 and 5 the extreme stimuli (Figure 2B). Specifically, per object class, an individual subject was exposed to either Items 1–5 or Items 3–7. Item 5 would be an extreme stimulus for a subject exposed to Items 1–5 (and the results associated with this stimulus in this subject would fall in Condition 5), whereas the same Item 5 would be a prototype for a subject exposed to Items 3–7 (and the results associated with

this stimulus in this subject would fall in Condition 3). In contrast to the design followed by all previous studies of norm-based encoding, this design assures that physical stimulus differences or preexisting stimulus biases cannot account for any differences in neural activity between the different conditions.

In our study, the prototype is defined as being the average stimulus of the stimulus distribution shown to a subject in the perceived stimulus space, which corresponds closely to the central stimulus in the parametric space as well as in the physical pixel-based stimulus space. This is often the case in studies of norm-based coding, but some exceptions exist, especially when the parametric space and perceived space do not fully correspond (Kayaert, Biederman, Op de Beeck, et al., 2005).

Behavioral Asymmetries during a Rating Task

To test for the development of biasing effects of the prototype, participants were asked to rate the similarity between pairs of exemplars from the same class in a block of 25 trials (all possible pairings of the five selected exemplars of a class). Each training block could be initiated by the subject and started with a preview of the five exemplars of a class, each shown three times, one stimulus at a time, in a random order (300 msec stimulus presentation, 300 msec ISI). Next, subjects could start the 25 trials. In each trial, two exemplars were shown sequentially (300 msec stimulus duration, 300 msec ISI), and subjects had to rate the dissimilarity between the two exemplars on a 7-point scale (1 = *same exemplar repeated*, 7 = *largest difference*) by typing in a number between 1 and 7 (Figure 3A). A small red fixation dot was present continuously.

The first training session took place on Day 1, included 16 blocks (four repetitions of the four classes in the fixed order and with the subject-specific group assignments), and lasted half an hour. The second training session took place on Day 2 and was an exact copy of the first session. The third training session took place on Day 3, right before scanning, and included four blocks (one repetition of each class in the subject-specific order).

Scanning

General

Scanning was performed posttraining (Day 3) at the Department of Radiology of the University Hospital in a 3-T Philips Intera magnet with an eight-channel SENSE head coil. The functional runs consisted of eight experimental runs and two localizer runs. The data of one experimental run of two subjects were removed because of excessive head motion (leaving 238 experimental runs). The localizer runs were acquired after the experimental runs and were followed by the acquisition of an anatomical scan.

For the experimental runs, functional images were acquired with an ascending EPI sequence (160 time points

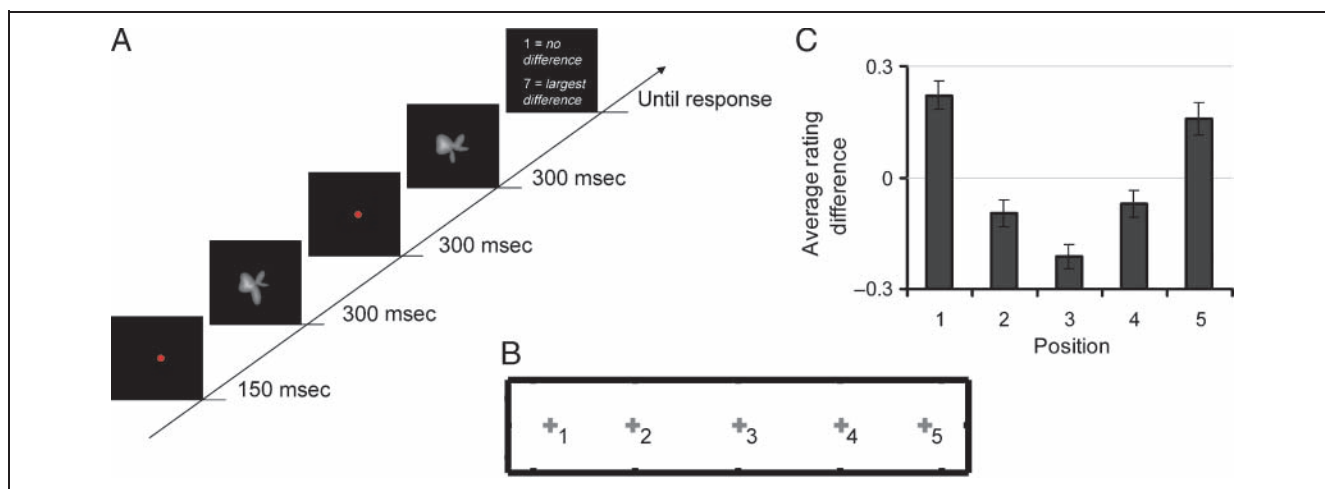


Figure 3. Behavioral training and data. (A) Trial structure during the similarity rating task. (B) Spatial configuration of the perceived similarities averaged across the four shape spaces. This configuration is derived from the similarity ratings using metric MDS. The spacing between the conditions is chosen to optimally reflect the rated difference. This configuration is very similar in dimensionality, stimulus order, and relative stimulus differences to the position of the shapes in parametric space (as illustrated in Figure 2). (C) The difference (asymmetry) in the rated similarity between trials with the same stimuli presented in a different order. For trials containing the prototype (Position 3), the rated similarity was higher when the prototype was presented second compared with when it was presented first. For trials containing extremes (Position 1 or 5), the rated similarity was higher when the extreme was presented first compared with when it was presented second. Error bars represent SEM across subjects.

per time series; repetition time = 2 sec, echo time = 30 msec, flip angle = 90° , voxel size = 2.75×2.75 mm, 36 slices, including most of the cortex except the frontal cortex and most superior parts of the parietal cortex; slice thickness = 3 mm and no interslice gap). The protocol was slightly different for the localizer runs (105 time points per time series; repetition time = 3 sec). A T1-weighted anatomical image (resolution $1 \times 1 \times 1.2$ mm) was acquired at the end of each session.

Experimental Runs

Eight experimental runs were collected per subject (two repetitions of the four classes in the fixed order and with the subject-specific group assignments; see Experimental design section). Each of the eight event-related experimental runs consisted of 160 trials. Each run started and ended with five fixation trials and included 20 presentations of each of the five exemplars of a class, 40 fixation trials, and 10 catch trials (see below). Before each run, we presented the same preview as during training (each exemplar presented three times in a random order).

Each trial lasted 2000 msec and contained one image. This image was presented for 300 msec, followed by an ISI of 1700 msec. A small red fixation dot was present throughout the run (Figure 5A). The order of the seven trial types (fixation, five exemplar positions, and catch) in each run was determined a priori by generating eight orders with preoptimized first-order counterbalancing with optseq2 software (NMR Center; Massachusetts General Hospital, Boston, MA). Per pair of participants, one received these orders in ascending sequence (1 to 8) and the other in descending sequence (8 to 1).

Subjects performed a luminance change detection task. They were asked to press one button whenever they saw a stimulus that was brighter than usual (catch trials). The stimulus shape was chosen randomly from the five stimuli in that block. The brightness difference between catch trials and normal trials in the first run was high and therefore the task easy. If detection performance was higher than 90% in a run, then the brightness difference was decreased in the next run to keep the task interesting and nontrivial. The catch trials were not included in the fMRI analysis.

Localizer Runs

The two localizer runs were used to define the object-selective LOC and its subregions (lateral occipital [LO], posterior fusiform [PFS]). Each localizer run consisted of 21 blocks and lasted 315 sec. In each block, we presented 20 images of the following: only a centrally presented fixation spot (five blocks per run), intact object images (eight blocks; 20 gray-scale object images were used of, e.g., flower, guitar, and tomato), or Fourier phase-scrambled versions of these object images (eight blocks). Each block lasted 15 sec. Intact and scrambled object images were presented in one of three colors (green, red, and blue). Subjects performed a color-change detection task. They were asked to press a button each time the stimuli changed color. This occurred three times per block.

Analysis of Imaging Data

Data were analyzed using the Statistical Parametric Mapping software package (SPM5; Wellcome Department of

Cognitive Neurology, London) as well as custom Matlab code.

Preprocessing

Images were first corrected for differences in acquisition time (slice timing; first slice as reference slice). They were then realigned to correct for head movements. The functional images of each subject were coregistered with his or her anatomical image, segmented and spatially normalized to a Montreal Neurological Institute template. During the normalization, the functional images were resampled to a voxel size of $3 \times 3 \times 3$ mm. The localizer data were spatially smoothed using a 6-mm FWHM kernel. For the ROI analysis, no smoothing was applied to the data from the experimental runs. For the whole-brain analysis, the functional data were spatially smoothed using an 8-mm FWHM kernel.

fMRI Data Analysis

LOC Analysis

For the localizer scans, the BOLD signal was modeled using a boxcar response model smoothed with a hemodynamic response function (Friston, 2003). The general linear model contained two independent variables (intact and scrambled) and six regressors (the translation and rotation parameters from the realignment) for each of both runs. Object-selective areas were defined as those areas in the LO and ventral occipito-temporal cortex (PFS) with a larger signal while subjects viewed intact objects versus textures. The *t* map corresponding to that contrast (thresholded at an uncorrected *p* value of .000001) was overlaid on the individual's coregistered anatomical image to select the voxels in the areas of interest using custom Matlab code. Figure 5B illustrates the LOC for one representative subject, including the subregions LO and PFS. The average number of voxels in left LO, right LO, left PFS, and right PFS was 883, 872, 535, and 585, respectively.

ROI Analysis

The time course of the BOLD signal intensity from the experimental runs was extracted by averaging the data from all the voxels within the independently defined ROIs. We averaged the signal intensity across the trials in each condition from -4 to 16 sec post-trial onset (11 time points). These event-related time courses of BOLD signal intensity were then converted to percent signal change (PSC) by subtracting the corresponding values for the fixation condition and dividing by that value (see Figure 6A and B). We then computed the peak for the time courses across conditions for each subject (on average occurring 4 sec after trial onset). The PSC at the peak served as the measured response for each condition and was entered in a repeated measures ANOVA (Figure 6C). For ease of comparison

with the behavioral data in Figure 3C, we also calculated a normalized PSC by dividing each peak response by the peak response for Position 3, for each subregion separately (Figure 6D).

Whole-brain Analysis

The experimental design was optimized for an ROI approach in which the peak response at a specific time point could be extracted, thus without any reliance on assumptions about hemodynamic response functions (Jiang et al., 2007; Yovel & Kanwisher, 2004; Kourtzi & Kanwisher, 2000, 2001). The design is less suited for whole-brain analyses that rely on convolution with continuous response functions, for which it would have been more optimal to add a random offset to the timing of each trial. Nevertheless, we performed a whole-brain analysis to verify that our ROI-based analyses would not miss important effects in other brain regions (note that superior parietal and frontal cortex were not included in the imaged volume). For each subject, we first performed a first-level whole-brain analysis using an informed basis set and eight conditions (five positions, catch trials, fixation trials, and fixation at the start and end) and six regressors (the translation and the rotation parameters from the realignment) per experimental run. Two contrasts were defined: objects versus fixation trials and extremes (Positions 1 and 5) versus prototype (Position 3). Next, second-level analyses were performed to make inferences about these contrasts at the population level.

RESULTS

Behavioral Evidence for a Norm-based Encoding

Several previous studies have proposed that asymmetries in a successive stimulus comparison task can be used as an index that some stimuli have a special status in the representation of stimulus differences (Op de Beeck et al., 2003a; Nosofsky, 1991). In each trial of the behavioral task, subjects rated the difference of two exemplars from the same class (Figure 3A). As shown in Table 1A, for each subject and position *X* (1 to 5) within each group, we first calculated the average difference rating (across classes and repetitions), once across those trials in which the exemplar at condition *X* was the first stimulus (SR1) and once across those trials in which the exemplar at position *X* was the second stimulus (SR2; "same" trials with twice the same stimulus were not used).

First, the difference in average rating was calculated by subtracting SR1 from SR2 (Table 1, last column, or Figure 3C). Figure 3C shows, on the basis of all 36 training blocks per subject, that there was a reliable asymmetry in ratings. For trials containing prototypical exemplars (Condition 3), the average rated difference was higher when the prototypical stimulus occupied the first position compared with when it occupied the second position in a

Table 1. Behavioral Data

Position	SR2	SR1	Average (SR1 and SR2)	SR2 – SR1
<i>A</i>				
1	4.68	4.46	4.57	0.22
2	3.88	3.87	3.88	–0.10
3	3.42	3.63	3.53	–0.21
4	3.75	3.82	3.79	–0.07
5	4.64	4.48	4.56	0.16
<i>B</i>				
1	3.52	3.23	3.37	0.28
2	3.25	3.29	3.27	–0.03
3	3.42	3.63	3.53	–0.21
4	3.23	3.23	3.23	0.00
5	3.44	3.25	3.35	0.19

Average dissimilarity values for each position (1 to 5) and stimulus order (SR1 and SR2). (A) All trials. (B) Trials with two steps or less between both stimuli (no trials with stimuli occupying Positions 1 and 5, 1 and 4, or 2 and 5).

trial. For trials containing extremes (Condition 1 or 5), the average difference was higher when these extremes occupied the second position in a trial compared with the first position. A repeated measures ANOVA revealed a significant difference in asymmetry between the five stimulus

conditions, $F(4, 116) = 18.81, p < .0001$. Planned comparisons showed a significant difference in asymmetry between stimulus Conditions 3 and 5, $F(1, 29) = 36.12, p < .0001$, and between stimulus Condition 3 and both extreme Conditions 1 and 5, $F(1, 29) = 60.04, p < .0001$.

Second, apart from the systematic pattern in terms of asymmetry, there is also a clear difference in the mean dissimilarity between the prototype Condition 3 and the extreme Conditions 1 and 5, with a smaller rated difference for the prototypes (Table 1A). This is a consequence of the fact that for extreme conditions, there are some stimulus conditions with a very large perceived difference, for example between stimuli at Positions 1 and 5. If we restrict the calculations of mean and asymmetry to the parametric steps that are present for the prototype (so a maximum of two steps), then the mean dissimilarity is about the same across all conditions, but the asymmetry is still very large, as shown in Table 1B.

Third, these asymmetries were present in each of the three training sessions (Figure 4). Interestingly, the effect size (approximately 0.4-unit difference between SR2 and SR1 for Position 3 and the average SR2–SR1 of Positions 1 and 5) did not change across sessions.

Finally, we verified that there was a close resemblance between the parametric and the physical shape manipulations and perceived similarity (see also Methods section). We constructed the perceived dissimilarity matrix as derived from the behavioral ratings (Table 2B). The perceived differences conform very closely to what one would expect on the basis of a one-dimensional parametric variation with constant steps. In fact, the correlation between the

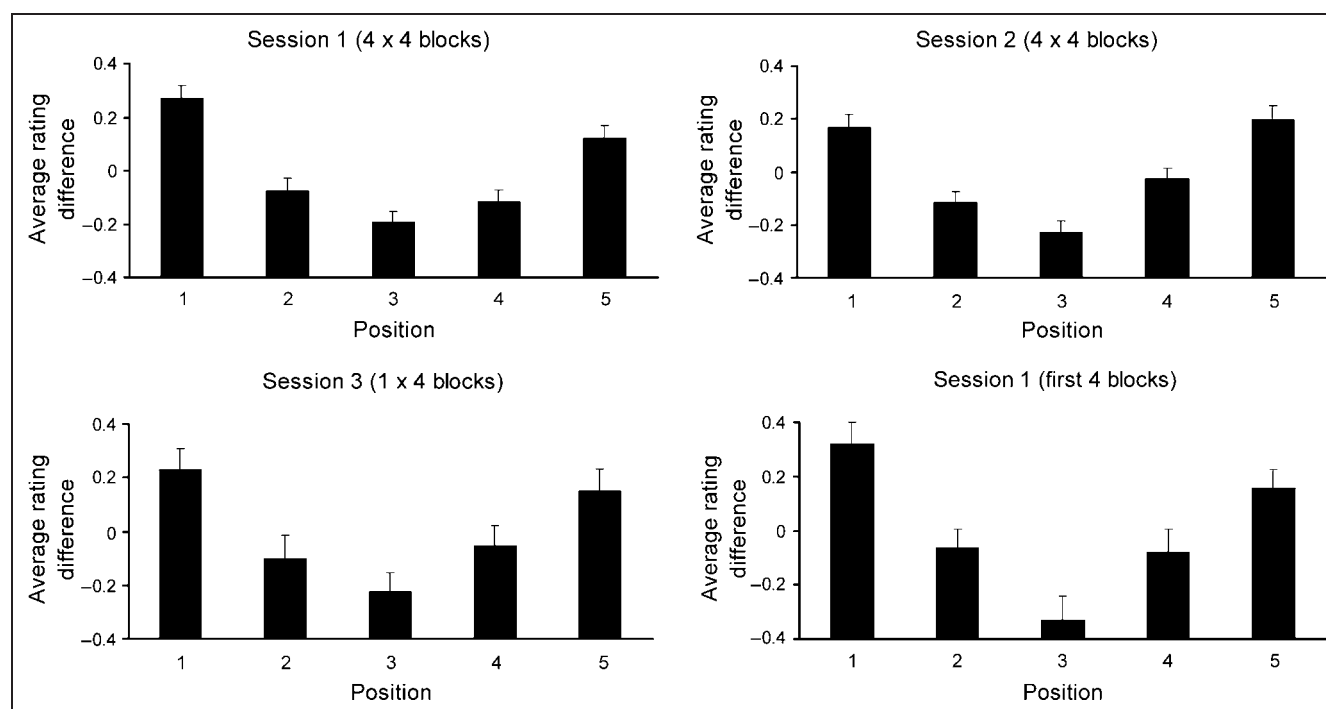
**Figure 4.** Behavioral asymmetries for separate sessions. Conventions are the same as in Figure 3C.

Table 2. Dissimilarity Matrices

	A					B				
	1	2	3	4	5	1	2	3	4	5
1	0	0.32	0.45	0.55	0.63	1.58	2.43	4.04	5.31	6.04
2	0.32	0	0.32	0.45	0.55	2.58	1.65	2.88	4.40	5.64
3	0.45	0.32	0	0.32	0.45	4.45	2.91	1.68	2.79	4.37
4	0.55	0.45	0.32	0	0.32	5.60	4.41	2.74	1.65	2.52
5	0.63	0.55	0.45	0.32	0	6.08	5.36	4.02	2.49	1.57

Rows represent SR1 (1 to 5) and columns SR2 (1 to 5) in a trial. (A) Pixel-wise differences averaged across the four actual stimulus spaces shown in Figure 2 in the same way as for the fMRI analyses, so for a subset of five images of all seven images in a space. Similar results were obtained with all images. (B) Perceived dissimilarity matrix from the similarity ratings (1 = *same*, 7 = *maximally different*) averaged across the four classes.

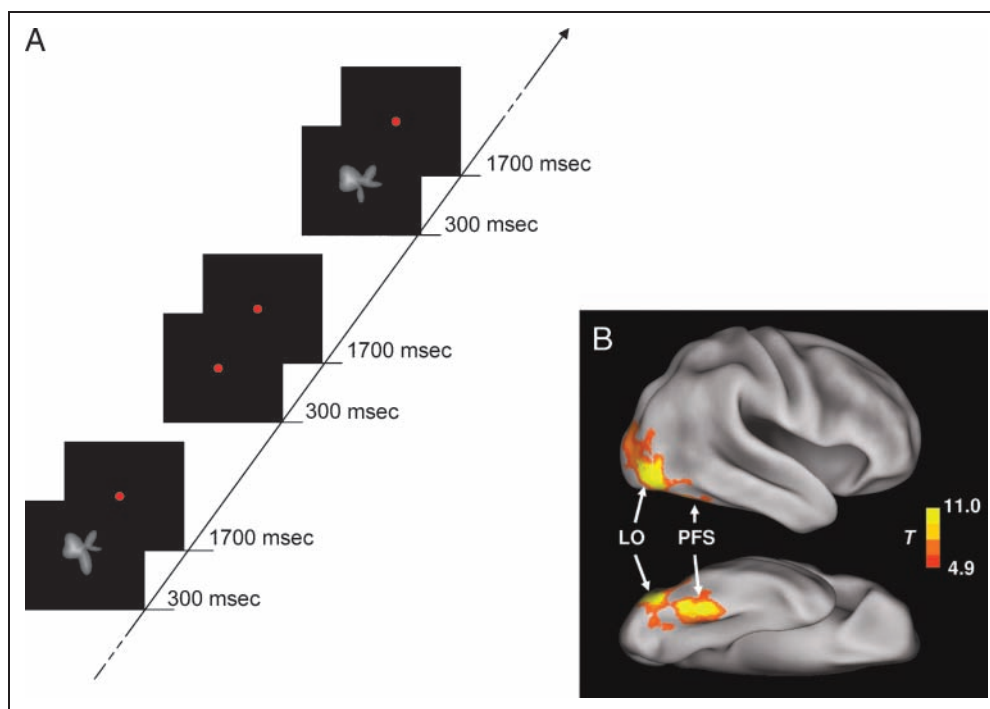
parametric differences and the perceived dissimilarity was very high ($r = .99$). Thus, the parametrically controlled shapes result in a good control of perceptual variability. This conclusion is confirmed by the application of metric multidimensional scaling (MDS) to the matrix with perceived dissimilarities (after making this matrix symmetric). MDS results in a low-dimensional spatial configuration of the five conditions in which the distance is as close as possible to the dissimilarity for each pair of conditions. A one-dimensional MDS solution explained 99% of the variance in the dissimilarity matrix and is very similar to what we would expect given the parametric shape manipulations. This spatial configuration is shown in Figure 3B. We can therefore be sure that the perceived average stimulus lies close to Position 3 and far from Positions 1 and 5.

In sum, we obtained behavioral evidence that a biasing effect from the perceived average stimulus has emerged from the first day in which the stimuli were viewed.

Neuroimaging Evidence for a Norm-based Encoding

To check the neural representation of the within-class stimulus distribution, we scanned the same subjects in an event-related fMRI experiment performed after the behavioral rating sessions (Figure 5A). Previous studies have indicated that objects are represented by neural activity in a large region in LO and ventral occipito-temporal cortex, referred to as the LOC (Grill-Spector & Malach, 2004; Kourtzi & Kanwisher, 2000; Malach et al., 1995). LOC is a

Figure 5. Main features of the design and main ROIs in the fMRI experiment. (A) Illustration of the sequence and timing of trials and images during part of a scan run. (B) The location of both subdivisions of LOC (LO and PFS) in one representative subject. The t map represents the contrast [intact object images–scrambled images] for one subject, thresholded at $p < 10^{-6}$ and shown on top of the PALS human atlas (top: lateral view; bottom: ventral view) using CARET software (Van Essen et al., 2001).



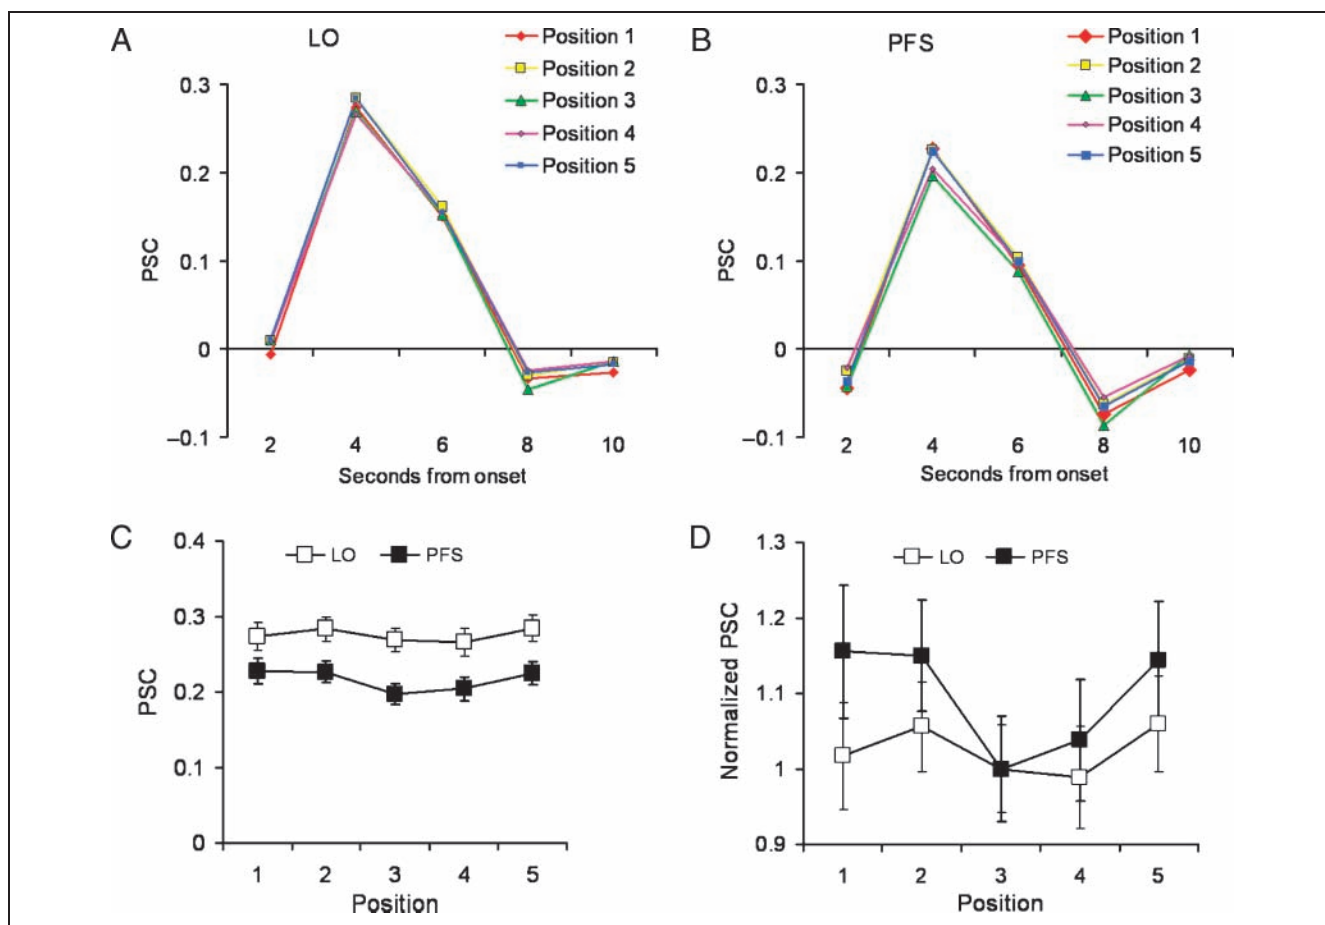


Figure 6. Effect of stimulus position in object-selective cortex. (A) PSC compared with the fixation baseline as a function of Position in LO. (B) PSC for each Position in PFS. (C) Peak responses as a function of ROI and Position. (D) Normalized PSC as a function of ROI and Position. Error bars represent *SEM* across subjects.

large brain region, and previous studies have suggested that it might be composed of two subregions (Grill-Spector, 2003; Grill-Spector et al., 1999): the LO gyrus (area LO) and the posterior fusiform region (area PFS). Area LO is generally thought of as being at a lower level in the object processing hierarchy compared with area PFS, giving us the opportunity to investigate how widespread the norm-based encoding is across the object vision pathway. We defined LO and PFS in each individual subject from their typical preference for intact object images over scrambled images (Figure 5B). We extracted the event-related fMRI response for each stimulus position in LO and PFS (Figure 6A and B), and we defined response strength at the peak of this event-related response (Figure 6C).

If visual experience can determine which stimulus is encoded as the prototype and which as an extreme class example, then we expect a lower neural population response for the prototype than for the extreme stimuli, although the attribution of physical stimuli to conditions was counterbalanced and reversed across subjects. A repeated measures ANOVA with factors Subregion (LO and PFS), Stimulus Rendering (two-dimensional silhouette or three-dimensional grassfire), and Position (1 to 5) on the

peak responses (Figure 6C) showed significant main effects of Subregion, $F(1, 28) = 20.23, p < .001$, and Position, $F(4, 112) = 3.16, p = .0167$. The main effect of Stimulus Rendering was not significant, $F(1, 28) = 0.85, p = .365$, and this factor did not interact with other effects. There was a significant interaction between Subregion and Position, $F(4, 112) = 2.99, p = .0219$. Thus, the variation in response strength across stimulus positions interacted significantly with ROI. Most importantly, the peak response (Figure 6C) was significantly lower for the prototype compared with the average of the two extreme conditions in area PFS, $F(1, 28) = 12.12, p = .0017$, but not in area LO, $F(1, 28) = 1.60, p = .217$. When comparing the response to extreme Position 5 with that to the prototype (so with stimuli perfectly counterbalanced across subjects), a significant difference is present in PFS, $F(1, 28) = 8.91, p = .0058$, but not in LO, $F(1, 28) = 2.53, p = .123$. The response to extreme Position 1 is also larger than that to the prototype in PFS, $F(1, 28) = 7.913, p = .0088$, but not in LO, $F(1, 28) = 0.23, p = .64$. Further division of PFS according to left versus right hemisphere revealed a significant difference between the prototype and the average response to the extremes in each

hemisphere (left PFS, $p = .004$; right PFS, $p = .003$). Also, the prototype effect in PFS tended to be graded instead of all or none, with stimuli at intermediate positions eliciting an intermediate response: The response to the prototype was also significantly smaller than the average response to the intermediate stimuli at Positions 2 and 4 ($p = .03$), and the response to Positions 2 and 4 tended to be smaller than the response to extremes (Positions 1 and 5, $p = .088$).

For ease of comparison with the behavioral data in Figure 3C and the previous report of norm-based encoding for faces in the monkey brain (Figure 4A in Leopold et al., 2006), we normalized these peak responses by dividing each by the average peak response for Position 3, separately for each subregion (Figure 6D). This representation reveals a similar U-shaped function centered around the prototype as found in Figure 3C, again only for area PFS. A repeated measures ANOVA on these normalized PSCs showed a significant main effect of Position, $F(4, 112) = 3.39$, $p = .0118$, and again only a significant interaction between Position and Subregion, $F(4, 112) = 4.133$, $p = .0037$. There were no main effects of Subregion, $F(1, 28) = 1.68$, $p = .21$, nor Stimulus Rendering, $F(1, 28) = 0.75$, $p = .39$. Planned comparisons showed the same significant differences as above, such as a lower normalized response to the prototype (Position 3) compared with the average response to the extremes (Positions 1 and 5) in area PFS, $F(1, 28) = 12.12$, $p = .0017$, but not in area LO, $F(1, 28) = 1.60$, $p = .217$.

These results indicate that a norm-based encoding is strongly modulated specifically in a higher level region of the object vision pathway after visual experience with the distribution of exemplars. This visual experience allows the representations to be centered or “anchored” on the prototype or norm. The effect size is relatively small, with the response being approximately 16% larger to extremes than to prototypes, but previous monkey single-

cell studies with faces (Freiwald, Tsao, & Livingstone, 2009) and novel shapes (Kayaert, Biederman, Op de Beeck, et al., 2005) have also found effect sizes close to 20%.

We also performed a whole-brain second-level analysis to find out whether other brain regions outside LOC would also show this effect. Of course, a second-level analysis for the contrast “objects versus fixation” showed significant activity throughout the ventral object vision pathway, including the typical anatomical location of retinotopic areas V1–V4 and LO and fusiform gyrus ($p < .001$ corrected). However, when comparing the average response to the extremes with that to the prototype, the contrast “extremes versus prototype” showed no voxels surviving the threshold ($p < .05$ corrected). At an uncorrected threshold of .05, there were one cluster of active voxels in the right fusiform gyrus and one in lateral posterior parietal cortex bilaterally but more extensive in the right hemisphere. Clear event-related responses were observed in these two second-level ROIs as shown in Figure 7. Prefrontal and superior parietal and frontal cortices were not covered during scanning.

Further analyses on the defined PFS ROI suggested that this apparent lateralization in the whole-brain analysis (only activity in the right hemisphere), which was not found in the ROI analysis, was due to more overlap between subjects in the anatomical location of PFS in the right hemisphere compared with the left hemisphere. For each ROI and subject separately, we calculated the number of voxels that matched the same ROI in each of the 29 other subjects separately. Second, we calculated the average of these 29 numbers for each subject and ROI. A paired t test across subjects showed that on average, the number of matching voxels across subjects is higher in right PFS ($n = 193$) compared with left PFS ($n = 170$), $t(29) = 2.37$, $p = .0247$. Thus, the apparent lateralization in the whole-brain analysis is probably due

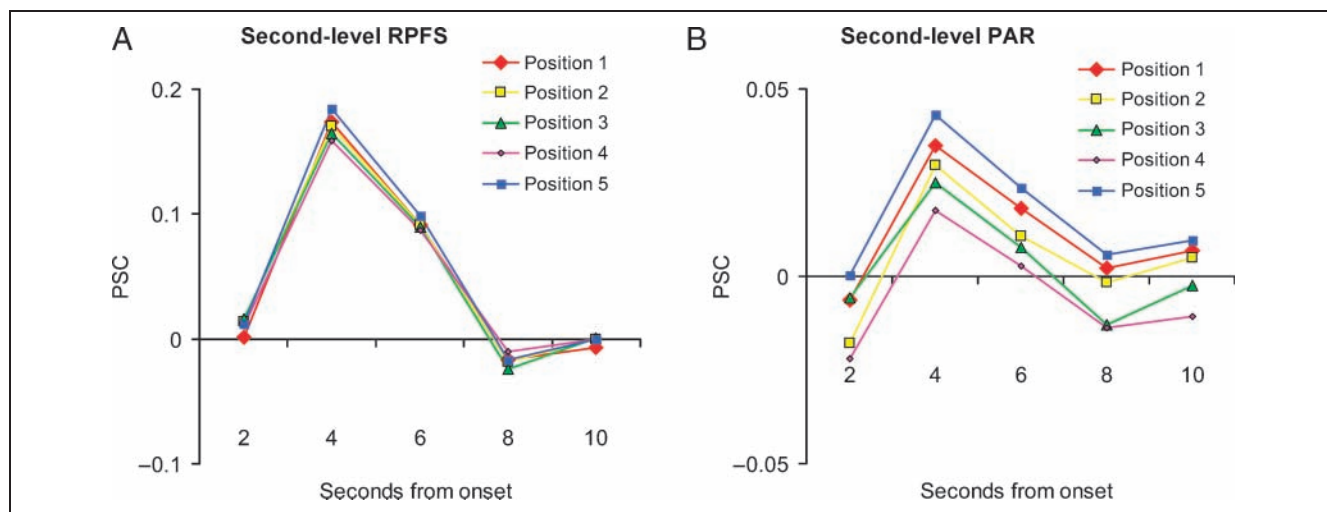


Figure 7. Second-level ROIs analysis. (A) PSC compared with the fixation baseline as a function of Position in the right posterior fusiform (RPFS) ROI identified by the second-level analysis (uncorrected $p < .05$). (B) PSC in the bilateral parietal (PAR) ROI identified in the same second-level analysis.

to between-subject variation in functional anatomy, which does not affect ROI analyses with ROIs defined in individual subjects.

Finally, we investigated whether the lower population response for prototypes can be explained by short-term adaptation (Grill-Spector & Malach, 2001; Kourtzi & Kanwisher, 2000), as has been suggested by N. Davidenko, D. Remus, and K. Grill-Spector (*Society for Neuroscience Abstracts*, 19.5, 2008) for the case of a block design. Short-term adaptation refers to a phenomenon in which the response to a stimulus depends on the preceding stimulus. The effect is modulated by the similarity between stimuli (more adaptation with higher similarity; Panis, Vangeneugden, Op de Beeck, & Wagemans, 2008), and it declines as a function of the time interval between adapter and test stimulus. Short-term adaptation between successive trials might indeed explain part of the lower responses for prototypical images found in the literature.

In our event-related study, the exact stimuli preceding each condition (one-back, i.e., the previous trial) were equated to allow an unbiased measure of the response in each condition. As a consequence, however, on average the prototype will be more similar to its preceding stimulus than an extreme will be just because an extreme stimulus might sometimes be preceded by a very different stimulus (the other extreme). We designed the experiment to avoid short-term adaptation as much as possible, more than most of the studies in the literature: We used an event-related design, each stimulus was only presented shortly, and there was a 1700-msec interval between successive images. Nevertheless, short-term adaptation cannot simply be excluded as a possibility.

We performed an explicit test of the presence of short-term adaptation between successive trials. A finer division of conditions on the basis of what occurred in the previous trial is not an acceptable approach for most conditions, given that this procedure would typically mess up the careful one-back randomization of the preceding stimulus referred to earlier. However, a powerful test is available from trials with extreme stimuli after extreme stimuli (Figure 8). In some cases (“same extreme” or SE), an extreme stimulus will succeed itself, resulting in maximal adaptation. In other cases (“different extreme” or DE), an extreme stimulus will succeed the other extreme stimulus, resulting in the maximal release of adaptation achievable in our design. Importantly, these two trial types are perfectly counterbalanced with regard to the physical stimuli. No significant difference could be detected between SE and DE in PFS (two-tailed paired t test), $t(29) = -.30$, $p = .77$, nor in LO (two-tailed paired t test), $t(29) = -.11$, $p = .91$, arguing against adaptation between successive trials in this experiment (a similar test with the same outcome was performed by Freiwald et al., 2009).

This lack of trial-to-trial adaptation stands in contrast to how these two “extreme-after-extreme” conditions compare with a third condition in which a prototype was preceded by an extreme stimulus (condition EP; this is a

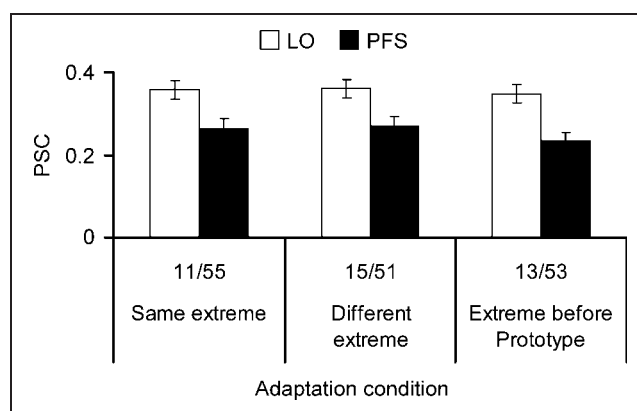


Figure 8. Test of trial-to-trial adaptation. The response is shown for three conditions: extreme stimulus after the same extreme stimulus (Position 1 after 1 or 5 after 5), extreme stimulus after the other extreme stimulus (Position 1 after 5 or 5 after 1), and prototype stimulus (Position 3) after extreme stimulus (1 or 5). Error bars represent *SEM* across subjects.

small subset of the trials in Condition 3 in the main analyses). So here the preceding stimulus is counterbalanced (always extreme), and the current stimulus is a prototype. As expected and consistent with the ROI analysis results, the response was lower in condition EP compared with the average of SE and DE in PFS (one-tailed paired t test), $t(29) = -1.85$, $p = .0376$, but not in LO (one-tailed paired t test), $t(29) = -.689$, $p = .248$.

DISCUSSION

The present study provides the first experimental evidence that empirical effects typically associated with norm-based encoding can be strongly modulated, even reversed among stimuli, after learning at a time scale that is longer than single successive trials but far shorter than the weeks of testing involved in previous neurophysiological studies of norm-based encoding with faces.

We found behavioral and neural evidence for a norm-based encoding scheme with novel object classes, which was restricted to a higher level region of the ventral visual system in the fusiform gyrus. We can conclude that this result is a result of visual experience with the distribution of exemplars from the object classes and not of preexisting stimulus biases because the critical difference between our design and previous studies of norm-referenced coding with faces and objects is that we changed the stimulus distribution between subjects, which allowed us to dissociate between the actual stimuli and the prototypicality status of a stimulus in the experienced distribution of stimuli.

Also in contrast to all previous fMRI studies on this topic, we used an event-related design with one image per trial and minimized trial-to-trial adaptation by having a relatively long time interval between successive images. Given

that we also verified empirically that no adaptation occurred specifically between successive images, we can conclude that the evidence for norm-based encoding relates to a more long-term effect that integrates across multiple trials.

Apart from revealing the role of visual experience for norm-based encoding, our findings show that this encoding scheme is limited within the ventral visual pathway to the part of object-selective cortex around the fusiform gyrus. No norm-based encoding was found in the LO gyrus. The latter region is typically considered as lower in the hierarchy of visual processing levels, and it has more sensitivity for low-level features such as retinal stimulus position (Sayres & Grill-Spector, 2008; Hemon, Kanwisher, & Op de Beeck, 2007). Nevertheless, it also codes for high-level aspects of stimuli that are represented in the fusiform region, including object shape (Haushofer, Livingstone, & Kanwisher, 2008; Op de Beeck et al., 2008; Kourtzi & Kanwisher, 2000). Thus, area LO appears to encode the shape features of both novel and more familiar objects without any reference to the class prototype. Note that norm-based encoding (see Figure 1) is a different type of representation than the so-called prototype-based models in the category learning literature (e.g., Ashby & Maddox, 1993). In fact, norm-based encoding integrates exemplar-based and prototype-based models because single neurons are tuned for class exemplars (Freiwald et al., 2009; Leopold et al., 2006; Kayaert, Biederman, Op de Beeck, et al., 2005), but there is an unequal distribution of the exemplars preferred by single neurons, as most neurons will prefer an extreme exemplar (Leopold et al., 2006).

The changes underlying norm-based encoding occur at a time scale that is similar to the time scale underlying other behavioral and neural phenomena such as long-term priming (Grill-Spector, 2001), which works across minutes and hours (Vuilleumier, Henson, Driver, & Dolan, 2002; Wiggs & Martin, 1998). Our behavioral data suggest that norm-based encoding does not need multiple days to develop, given that the behavioral asymmetry effect was already found in the first 10 minutes or 100 trials (i.e., after the first four blocks of the first behavioral session, see Figure 4). We performed the fMRI sessions after multiple training sessions, but the positive result that we obtained in area PFS might not take multiple days to develop, and it might be working over tens of seconds or minutes. Thus, our experiment shows that visual experience can influence which stimulus is processed as the prototype and that this effect is building up across intervening trials, but it does not answer all questions about the exact time scale (minutes, hours, or days) over which these experience-related dynamics occur—further experiments will need to address this issue.

Interestingly, the parietal cortex—classically related with attentional processing (for reviews, see Corbetta & Shulman, 2002; Kanwisher & Wojciulik, 2000)—also tended to respond stronger to extreme compared with prototypical exemplars (Figure 7B). Although this result was not strongly

significant in a statistical sense, it is consistent with the idea that norm-based encoding has attentional advantages (Barlow, 1990), possibly by influencing the saliency of stimuli in a top-down fashion (see Corbetta & Shulman, 2002). If correct, we would predict that extreme exemplars pop out from a field of prototypical exemplars in a visual search task (preliminary evidence for this has been obtained by Kayaert, Panis, Op de Beeck, & Wagemans, 2009).

In our study, the prototype was conceptualized as the average. This is a common procedure, also widely applied in the studies on norm-based encoding with faces. Furthermore, it is also the classic way to define prototypes in the extended literature on prototype-based category learning (Smith & Minda, 1998; Posner & Keele, 1968). In most of these cases, as in our experiment, prototypicality is dissociated from raw frequency because the latter is either the same for prototypes and nonprototypes (as in our experiment) or even smaller for prototypes (e.g., in the classic prototype learning experiments the prototype is never shown during training and has to be “abstracted”). In everyday situations, raw frequency might be correlated with prototypicality and/or prototypicality might be dissociated from the average.

On a related note, stimuli might acquire a “special” status in the representation of other stimuli for other reasons than being a prototype (Nosofsky, 1991). First of all, it is not entirely clear which is the “special” stimulus on the basis of the neural data: Is it the stimulus associated with the smallest response or the one associated with the largest response? Studies on norm-based encoding interpret the smaller response for the “average” stimulus as evidence that this stimulus is the anchor used to compare all other stimuli with. However, all our data are showing is that stimuli get a different status in the mental space of category exemplars for representing the stimuli, and by itself, this does not prove which stimulus is the most important one. According to the adaptive coding model of Rhodes and Jeffery (2006) shown in Figure 1, however, the question of which stimulus is most important is not a good one because it refers to two sides of the same coin: Although neurons stay tuned to the (low or high) extreme values encountered in the natural world, they adapt to prevalent (average) values. Stimuli with average values on dimensions used to discriminate stimuli are only coded implicitly by equal activation of two neural populations, one adapted to respond strongest to above-average values and the other one to below-average values. As a result, neurons fire strongest to extreme stimuli, but stimuli with average values on many dimensions used to discriminate exemplars (i.e., stimuli close to the prototype) are represented with the lowest overall activity and are processed most efficiently. The prototype changes dynamically depending on the frequency with which exemplars are encountered, the identity of the exemplars experienced, the task factors, and so forth. Such dynamic changes are revealed at the neural level in our experiment.

Furthermore, we leave open the possibility that other factors besides the prototypicality of a stimulus might contribute to its function as an anchor. Some of these factors might not be related to the bottom-up statistics of the visual input. For example, a system might learn that specific stimuli are particularly important through feedback learning, and in this case these stimuli might function as anchors. A more general term that is appropriate for all these cases would be “anchor” instead of norm/prototype. Thus, using a more general terminology, we have revealed that visual experience with novel exemplars can cause some stimuli to get a special status as an anchor in the encoding of the differences between exemplars of a particular class. These flexible and experience-based anchoring mechanisms might be a universal type of representation in the brain, in vision and beyond (Chapman & Johnson, 2002; Tourangeau & Rasinski, 1988; Couclelis, Golledge, Gale, & Tobler, 1987).

Given the importance of the concept of norm-based encoding, we anticipate the presented findings to be a starting point for future studies looking into the role of several other variables for how this encoding scheme is affected by visual experience. One potentially important factor is the behavioral state of the subjects. In our experiments, subjects attended to the shape differences during the behavioral rating sessions. Several aspects of this training task might matter. First, subjects are trained in this discrimination rating task, providing the opportunity to refine perceptual representations. Second, the training task might make the shape differences more salient to subjects. Third, the “average” rating in trials including a prototype was lower than the rating in trials including extreme stimuli (see Table 1A, columns 2 and 3). So, on average, the extremes and the prototype are not associated with the same behavioral response during training. The same phenomenon actually occurs with familiar objects, for example, more average faces are rated as more attractive (Rhodes, 2006), and social reactions differ as a function of attractiveness.

During scanning, the task was very different: Subjects were required to press a button every time an object was brighter than the other objects, and shape differences were no longer relevant (these catch trials were excluded from the analyses). The strength of the norm-based encoding might of course depend on the relevance of shape differences before and/or during scanning (see, e.g., Gillebert, Op de Beeck, Panis, & Wagemans, 2009). Note that studies showing norm-based encoding (which did not take into account a possible role of within-experiment experience) have also mostly ignored the relevance of task context, for example, the two monkeys in the study of Leopold et al. (2006) were very different in this respect. More studies are needed to pinpoint the role of task context for norm-based encoding.

In conclusion, we found behavioral and neural evidence for a norm-based encoding scheme with novel object classes, which was restricted to a higher level region

of the ventral visual system in the fusiform gyrus. Thanks to the careful counterbalancing of physical stimulus properties across conditions, we can conclude that this finding is a result of visual experience with the distribution of exemplars from the object classes. Thus, not all exemplars of novel object classes are treated equally for the representation of within-class object differences. The distribution of exemplars can cause some stimuli to get a special status as an anchor in the encoding of the differences between exemplars of a particular class.

Acknowledgments

The authors thank Ron Peeters for technical assistance with the experiments, Greet Kayaert and Rufin Vogels for helpful discussions, and two anonymous reviewers for important comments.

This work was supported by a federal research action (IUAP P6/29), the Research Council of K.U. Leuven (IMPH/06/GHW), the Fund for Scientific Research—Flanders (G.0281.06 and 1.5.022.08), a Methusalem grant (METH/08/02) from the Flemish Government, and the Human Frontier Science Program (CDA 0040/2008).

Reprint requests should be sent to Hans P. Op de Beeck, Laboratory of Biological Psychology, University of Leuven (K.U. Leuven), Tiensestraat 102, B3000 Leuven, Belgium, or via e-mail: hans.opdebeeck@psy.kuleuven.be.

REFERENCES

- Anderson, N., & Wilson, H. (2005). The nature of synthetic face adaptation. *Vision Research*, *45*, 1815–1828.
- Ashby, F. G., & Maddox, W. T. (1993). Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology*, *37*, 372–400.
- Barlow, H. B. (1990). A theory about the functional role and synaptic mechanism of visual after-effects. In C. Blakemore (Ed.), *Vision: Coding and efficiency*. Cambridge: Cambridge University Press.
- Blum, H. (1973). Biological shape and visual science. *Journal of Theoretical Biology*, *38*, 205–287.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.
- Chapman, G. B., & Johnson, E. J. (2002). Incorporating the irrelevant: Anchors in judgments of belief and value. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases* (pp. 120–138). Cambridge: Cambridge University Press.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*, 201–215.
- Couclelis, H., Golledge, R. G., Gale, N., & Tobler, W. (1987). Exploring the anchor-point hypothesis of spatial cognition. *Journal of Environmental Psychology*, *7*, 99–122.
- De Baene, W., Premereur, E., & Vogels, R. (2007). Properties of shape tuning of macaque inferior temporal neurons examined using rapid serial visual presentation. *Journal of Neurophysiology*, *97*, 2900–2916.
- Drucker, D. M., & Aguirre, G. K. (2009). Different spatial scales of shape similarity representation in lateral and ventral LOC. *Cerebral Cortex*, *19*, 2269–2280.
- Freiwald, W. A., Tsao, D. Y., & Livingstone, M. S. (2009). A face feature space in the macaque temporal lobe. *Nature Neuroscience*, *12*, 1187–1196.

- Friston, K. (2003). Introduction: Experimental design and statistical parametric mapping. In R. S. J. Frackowiak, K. J. Friston, C. Frith, R. Dolan, C. J. Price, S. Zeki, et al. (Eds.), *Human brain function* (pp. 599–633). New York: Academic Press.
- Georgieva, S. S., Todd, J. T., Peeters, R., & Orban, G. A. (2008). The extraction of 3D shape from texture and shading in the human brain. *Cerebral Cortex*, *18*, 2416–2438.
- Gillebert, C. R., Op de Beeck, H. P., Panis, S., & Wagemans, J. (2009). Subordinate categorization enhances the neural selectivity in human object-selective cortex for fine shape differences. *Journal of Cognitive Neuroscience*, *21*, 1054–1064.
- Grill-Spector, K. (2001). Semantic versus perceptual priming in fusiform cortex. *Trends in Cognitive Sciences*, *5*, 227–228.
- Grill-Spector, K. (2003). The neural basis of object perception. *Current Opinion in Neurobiology*, *13*, 159–166.
- Grill-Spector, K., Kushnir, T., Edelman, S., Avidan, G., Itzhak, Y., & Malach, R. (1999). Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron*, *24*, 187–203.
- Grill-Spector, K., & Malach, R. (2001). fMR-adaptation: A tool for studying the functional properties of human cortical neurons. *Acta Psychologica*, *107*, 293–321.
- Grill-Spector, K., & Malach, R. (2004). The human visual cortex. *Annual Review of Neuroscience*, *27*, 649–677.
- Haushofer, J., Livingstone, M. S., & Kanwisher, N. (2008). Multivariate patterns in object-selective cortex dissociate perceptual and physical shape similarity. *PLoS Biology*, *6*, e187.
- Hemond, C. C., Kanwisher, N. G., & Op de Beeck, H. P. (2007). A preference for contralateral stimuli in human object- and face-selective cortex. *PLoS One*, *2*, e574.
- Jiang, X., Bradley, E., Rini, R. A., Zeffiro, T., Vanmeter, J., & Riesenhuber, M. (2007). Categorization training results in shape- and category-selective human neural plasticity. *Neuron*, *53*, 891–903.
- Kanwisher, N., & Wojciulik, E. (2000). Visual attention: Insights from brain imaging. *Nature Reviews Neuroscience*, *1*, 91–100.
- Kayaert, G., Biederman, I., Op de Beeck, H. P., & Vogels, R. (2005). Tuning for shape dimensions in macaque inferior temporal cortex. *European Journal of Neuroscience*, *22*, 212–224.
- Kayaert, G., Biederman, I., & Vogels, R. (2005). Representation of regular and irregular shapes in macaque inferotemporal cortex. *Cerebral Cortex*, *15*, 1308–1321.
- Kayaert, G., Panis, P., Op de Beeck, H. P., & Wagemans, J. (2009). *Atypical objects are easier to spot*. 32th European Conference on Visual Perception, Regensburg, Germany.
- Kourtzi, Z., & Kanwisher, N. (2000). Cortical regions involved in perceiving object shape. *Journal of Neuroscience*, *20*, 3310–3318.
- Kourtzi, Z., & Kanwisher, N. (2001). Representation of perceived object shape by the human lateral occipital complex. *Science*, *293*, 1506–1509.
- Lee, K., Byatt, G., & Rhodes, G. (2000). Caricature effects, distinctiveness, and identification: Testing the face-space framework. *Psychological Science*, *11*, 379–385.
- Leopold, D. A., Bondar, I. V., & Giese, M. A. (2006). Norm-based face encoding by single neurons in the monkey inferotemporal cortex. *Nature*, *442*, 572–575.
- Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, *4*, 89–94.
- Leopold, D. A., Rhodes, G., Muller, K. M., & Jeffery, L. (2005). The dynamics of visual adaptation to faces. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences*, *272*, 897–904.
- Loffler, G., Yourganov, G., Wilkinson, F., & Wilson, H. R. (2005). fMRI evidence for the neural representation of faces. *Nature Neuroscience*, *8*, 1386–1390.
- Logothetis, N. K., & Sheinberg, D. L. (1996). Visual object recognition. *Annual Review of Neuroscience*, *19*, 577–621.
- Malach, R., Reppas, J. B., Benson, R. R., Kwong, K. K., Jiang, H., Kennedy, W. A., et al. (1995). Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *92*, 8135–8139.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39–61.
- Nosofsky, R. M. (1991). Stimulus bias, asymmetric similarity, and classification. *Cognitive Psychology*, *23*, 94–140.
- Op de Beeck, H. P., Torfs, K., & Wagemans, J. (2008). Perceived shape similarity among unfamiliar objects and the organization of the human object vision pathway. *Journal of Neuroscience*, *28*, 10111–10123.
- Op de Beeck, H. P., Wagemans, J., & Vogels, R. (2001). Inferotemporal neurons represent low-dimensional configurations of parameterized shapes. *Nature Neuroscience*, *4*, 1244–1252.
- Op de Beeck, H. P., Wagemans, J., & Vogels, R. (2003a). Asymmetries in stimulus comparisons by monkey and man. *Current Biology*, *13*, 1803–1808.
- Op de Beeck, H. P., Wagemans, J., & Vogels, R. (2003b). The effect of category learning on the representation of shape: Dimensions can be biased but not differentiated. *Journal of Experimental Psychology: General*, *132*, 491–511.
- Panis, S., Vangeneugden, J., Op de Beeck, H. P., & Wagemans, J. (2008). The representation of subordinate shape similarity in human occipitotemporal cortex. *Journal of Vision*, *8*, 1–15.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353–363.
- Rhodes, G. (2006). The evolutionary psychology of facial beauty. *Annual Review of Psychology*, *57*, 199–226.
- Rhodes, G., Brennan, S., & Carey, S. (1987). Identification and ratings of caricatures: Implications for mental representations of faces. *Cognitive Psychology*, *19*, 473–497.
- Rhodes, G., & Jeffery, L. (2006). Adaptive norm-based coding of facial identity. *Vision Research*, *46*, 2977–2987.
- Rosch, E. (1973). On the internal structure of perceptual and semantic categories. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language*. New York: Academic Press.
- Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–439.
- Sayres, R., & Grill-Spector, K. (2008). Relating retinotopic and object-selective responses in human lateral occipital cortex. *Journal of Neurophysiology*, *100*, 249–267.

- Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 1411–1430.
- Suzuki, S., & Cavanagh, P. (1998). A shape-contrast effect for briefly presented stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1–27.
- Tourangeau, R., & Rasinski, K. A. (1988). Cognitive processes underlying context effects in attitude measurement. *Psychological Bulletin*, *103*, 299–314.
- Tsao, D. Y., & Freiwald, W. A. (2006). What's so special about the average face? *Trends in Cognitive Sciences*, *10*, 391–393.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327–352.
- Van Essen, D. C., Lewis, J. W., Drury, H. A., Hadjikhani, N., Tootell, R. B., Bakircioglu, M., et al. (2001). Mapping visual cortex in monkeys and humans using surface-based atlases. *Vision Research*, *41*, 1359–1378.
- Vuilleumier, P., Henson, R. N., Driver, J., & Dolan, R. J. (2002). Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nature Neuroscience*, *5*, 491–499.
- Wiggs, C. L., & Martin, A. (1998). Properties and mechanisms of perceptual priming. *Current Opinion in Neurobiology*, *8*, 227–233.
- Yovel, G., & Kanwisher, N. (2004). Face perception; domain specific, not process specific. *Neuron*, *44*, 889–898.