Cerebral Cortex Advance Access published November 21, 2007

Cerebral Cortex
doi:10.1093/cercor/bhm196

A Stable Topography of Selectivity for Unfamiliar Shape Classes in Monkey Inferior Temporal Cortex

Hans P. Op de Beeck[1,2], Jennifer A. Deutsch[1], Wim Vanduffel[3,4], Nancy G. Kanwisher[1,3] and James J. DiCarlo[1,3]

[1]McGovern Institute for Brain Research, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA, [2]Laboratory of Experimental Psychology, K.U. Leuven, B3000 Leuven, Belgium, [3]Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, MA 02129, USA and [4]Laboratorium voor Neuro- & Psychofysiologie, K.U. Leuven Medical School, Leuven B3001, Belgium

The inferior temporal (IT) cortex in monkeys plays a central role in visual object recognition and learning. Previous studies have observed patches in IT cortex with strong selectivity for highly familiar object classes (e.g., faces), but the principles behind this functional organization are largely unknown due to the many properties that distinguish different object classes. To unconfound shape from meaning and memory, we scanned monkeys with functional magnetic resonance imaging while they viewed classes of initially novel objects. Our data revealed a topography of selectivity for these novel object classes across IT cortex. We found that this selectivity topography was highly reproducible and remarkably stable across a 3-month interval during which monkeys were extensively trained to discriminate among exemplars within one of the object classes. Furthermore, this selectivity topography was largely unaffected by changes in behavioral task and object retinal position, both of which preserve shape. In contrast, it was strongly influenced by changes in object shape. The topography was partially related to, but not explained by, the previously described pattern of face selectivity. Together, these results suggest that IT cortex contains a large-scale map of shape that is largely independent of meaning, familiarity, and behavioral task.

Keywords: categorization, fMRI, learning, object recognition, primate, visual perception

Introduction

The ventral visual pathway in primates plays a central role in visual object recognition. This pathway culminates in inferior temporal (IT) cortex, a region whose functional organization is largely unknown (Tanigawa et al. 2005). Here we measured functional magnetic resonance imaging (fMRI) responses in monkey IT cortex to unfamiliar object classes to investigate the large-scale functional organization of this region.

fMRI studies in humans and monkeys have reported millimeter- to centimeter-scale patches of cortex in the ventral visual pathway that are selectively responsive to ethologically important classes of objects, such as faces, places, and body parts (Kanwisher et al. 1997; Epstein and Kanwisher 1998; Downing et al. 2001; Tsao et al. 2003; Pinsk et al. 2005). Furthermore, information about familiar object categories (e.g., chairs, scissors) is contained in the pattern of response across the ventral visual pathway (Haxby et al. 2001). These results suggest that IT cortex may be organized in terms of object category. However, most prior studies have measured fMRI responses only to familiar object categories. Thus, it is impossible to determine if the resulting functional organization reflects object shape, evolutionary significance, familiarity, meaning, type of processing, or associated eccentricity biases. Each of these properties has been proposed as an explanation for the specificity for object category found in the ventral visual pathway (Chao et al. 1999; Haxby et al. 2000; Kanwisher 2000; Tarr and Gauthier 2000; Beauchamp et al. 2002; Malach et al. 2002).

Given the ambiguity of cortical responses to familiar objects, our first goal was to determine if regions in monkey IT cortex respond differentially to novel object classes. Because novel objects do not obviously differ in evolutionary significance, meaning, and the type of processing they elicit, such objects might reveal little or no functional organization. Studies using techniques with higher spatial resolution, like optical imaging and extracellular recordings, have revealed a columnar organization in IT cortex (Fujita et al. 1992; Tsunoda et al. 2001; Tanigawa et al. 2005). However, with the exception of faces (Wang et al. 1998; Tsao et al. 2006), this organization is at a relatively small spatial scale. Thus, it is not clear whether any functional organization for unfamiliar object classes would be detectable with fMRI.

Our second goal was to understand how experience might alter the functional organization of IT cortex. For example, even if little functional organization was found for novel objects, then some organization might appear after large amounts of experience with one or more object classes. This hypothesis would be consistent with the finding from optical imaging work that only faces are associated with a functional organization that spans a larger area than the typical diameter of a column (around 0.5 mm) in IT cortex (Wang et al. 1998). Alternatively, if a clear functional organization was found, even for novel object classes, then this organization might be changed as some object classes became more familiar or behaviorally significant.

Thus, in considering what one might find using fMRI to examine the spatial organization of object selectivity in IT and the effect of object familiarity, there are 4 broad possibilities (Fig. 1): no topography of selectivity for novel or for familiar objects (Fig. 1*a*), no topography of selectivity for novel objects but clear selectivity for familiar objects (Fig. 1*b*), topography of selectivity for novel objects that does not change as objects become more familiar (Fig. 1*c*), and topography of selectivity for novel objects that changes as objects become more familiar (Fig. 1*d*).

To answer these questions, we carried out a series of contrast-enhanced fMRI studies in awake monkeys using initially novel object classes in which we set out to focus on monkey IT cortex. We found a clear topography of object-class selectivity for these novel objects across IT cortex. This selectivity topography was stable across months and across extensive training with one of the object classes. In a series of follow-up studies, we found that the selectivity topography was also stable across changes in behavioral task and across changes

in object position, but it was sensitive to changes in shape properties. The most parsimonious interpretation of these results is that adult IT cortex contains a large-scale (mm+) "map" of shape that is largely stable in adult animals.

Materials and Methods

Animals and Surgery

Experiments were performed on 2 rhesus monkeys (*Macaca mulatta*), 1 male, monkey J, and 1 female, monkey M (weights at start of experiments ~3.7 and ~3.2 kg, respectively). Before behavioral training, aseptic surgery was performed to attach an MR-compatible plastic head post to the skull of each monkey according to previously published methods (Vanduffel et al. 2001). During training, monkeys were seated in the so-called "sphinx" position in a plastic monkey chair. The experiments spanned 20 months for monkey J and 18 months for monkey M from this initial surgery to the last data acquisition. The monkeys were housed and trained in the behavioral tasks at the Massachusetts Institute of Technology and were transported to Massachusetts General Hospital for MRI. All animal procedures were performed in accordance with National Institutes of Health guidelines and were approved by the Massachusetts Institute of Technology Committee on Animal Care and the Massachusetts General Hospital's Subcommittee on Research Animal Care.

Eye Position Monitoring

All behavioral tasks and MR scanning involved control of eye position. Horizontal and vertical eye position was determined based on the pupil location using camera-based systems (mock scanner during training: EyeLink II with 250 Hz camera, SR Research, Osgoode, Ontario, Canada; MR scanner, Iscan Inc. with 120 Hz camera, Cambridge, MA). All further procedures for processing the eye data were similar to previous single-unit physiology experiments (DiCarlo and Maunsell 2000).

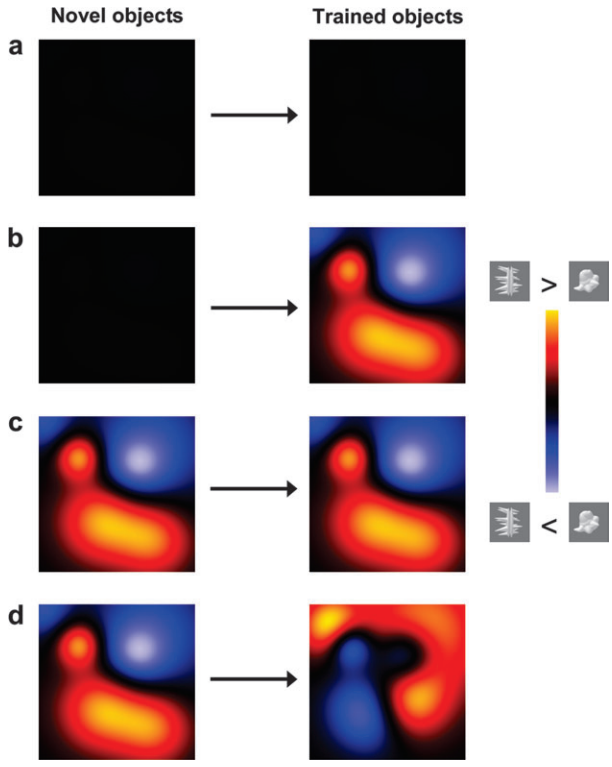

Figure 1. Four qualitative hypotheses about how much selectivity for initially novel object classes might be found with fMRI and how this selectivity might change as a consequence of training to gain expertise with those object classes. (*a*) No topography of selectivity before or after training; (*b*) no topography of selectivity before training but clear selectivity after training; (*c*) topography of selectivity before training that is unaffected by training; (*d*) topography of selectivity that changes as a consequence of training.

Visual Stimuli

Stimuli were presented on a video monitor (43 × 30 cm; 75 Hz frame rate; 1920 × 1200 pixels) positioned at 81 cm from the monkeys (mock scanner) or back projected from an LCD projector onto a screen (29 × 22 cm; 60 Hz frame rate; 1024 × 768 pixels) positioned at ~50 cm from the monkeys (MR scanner). Stimuli had approximately the same size (stimulus outline of 6 by 6 visual degrees) in the mock scanner and in the MR scanner.

Initial behavioral training of the monkeys in the behavioral tasks involved a broad range of stimuli (including natural scenes; pictures of animals, cars, and other everyday objects; human faces; and Fourier-scrambled images). The main set of experiments involved 3 parameterized classes of novel objects ("smoothies," "spikies," and "cubies," see Fig. 2). These classes were chosen to differ on "global" shape features such as aspect ratio and object structure (e.g., a horizontal base with vertical parts for the cubies), as well as on "local" shape features such as the presence of curved versus straight contours (smoothies vs. cubies) and the proportion of acute and obtuse angles (spikies vs. smoothies). The division and naming of these 2 sets of properties are somewhat arbitrary, but it reflects a clear distinction. For example, the manipulation of what we refer to as local shape features does not induce large changes of the global aspect ratio of the objects.

These stimuli were created with custom algorithms written in Matlab. Exemplars in each object class differ on 4 complex shape dimensions. Most dimensions were a composite of several simple shape parameters (e.g., size/thickness of several shape protrusions), and different dimensions changed aspects in different locations on the stimuli. Thus, objects could not be discriminated by looking at only a small part of each object, and the whole stimulus shape had to be taken into account to attain good discrimination performance. Each class contained 1296 exemplars.

One additional experiment included versions of these original 3 object classes in which the aforementioned global and local shape features were mixed ("spiky smoothies," cuby spikies, and smoothy cubies, see Fig. 9). The manipulation of global and local shape was a relative one, that is, the manipulation of one property automatically induced changes (albeit to a lesser extent) in the other property.

Localizer scans included a stimulus set composed of pictures of 20 rhesus monkey faces, pictures of 20 objects that were familiar to the monkeys (e.g., fruits and toys), and 20 Fourier-scrambled square images (with a Fourier spectrum that corresponded to the power spectrum of the 20 object pictures).

Behavioral Tasks

Each monkey was trained in 3 behavioral tasks that were learned successively using standard operant conditioning techniques (positive reinforcement with fluid while monkeys were under fluid restriction). Each monkey initially learned all tasks with a wide range of stimuli that did not include the 3 novel object classes used in the experiments. During the behavioral training, the monkeys were also acclimated to

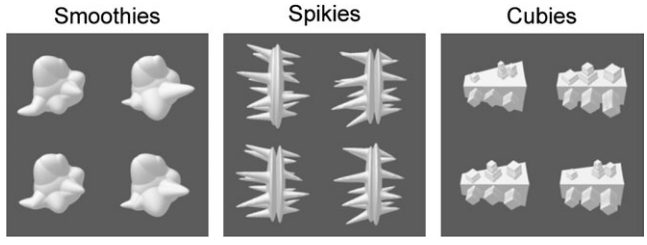

Figure 2. Four example objects from each of the 3 parameterized object classes used throughout most experiments (smoothies, spikes, cubies). In each panel, the 2 example objects in the top row are maximally distinct from each other (100%) within the space of possible objects in the class. For comparison, the 2 example objects in the bottom row are only 60% as distinct from each other.

factors associated with fMRI scanning, most importantly high acoustic background noise and intravenous (i.v.) injections prior to the session as necessary for the administration of the MRI contrast agent.

Passive Fixation

In this task, the monkey was intermittently rewarded for maintaining fixation within a square-shaped central fixation window (on average less than 2.5 visual degrees square). The interval between rewards was systematically decreased (never less than 1200 ms) as the monkey maintained its fixation within the window. The monkey learned to do this task with only a fixation spot (small square of 0.3 visual degrees on a side) and with stimuli appearing together with the fixation spot. The stimulus presentation was independent of the monkey's fixation behavior (i.e., stimuli were also presented when the monkey was not fixating).

Color Task

In this task, the monkey was also intermittently rewarded according to the passive fixation task requirements (above), but on top of that the monkey received additional rewards for correctly detecting changes in the color of successively presented object stimuli. Most objects had the same color (weakly saturated red, green, or blue) as the preceding object, but the color changed for a small proportion of the objects (not correlated with stimulus shape). After such a color change, the new color would be the same for several successive objects. At the end of the learning phase for this "color task" and during scanning with this task, there was 1 color change per 20 object stimuli presented. The monkey was trained to signal a detected color change by saccading to a small white square 8 degrees above the central fixation spot. The monkey was rewarded more for signaling a color change than for keeping fixation (up to 100% more). This proportion of reward for fixation and color change detection was important because we wanted high overall fixation percentage, high hit rate for stimuli with a changed color, and a low false alarm rate for stimuli with no change in color. The reward for correct responses was adapted to keep the false alarm rate below 5%. The monkey never had to switch abruptly from the passive fixation task to the color task, although the 2 tasks were sometimes performed on successive days. Each monkey accomplished this switch without apparent cost, and there are several task differences that could serve as a cue (the presence of the response corner, a white square, in the color task; colored vs. gray stimuli).

Shape Discrimination Task

In this task, the monkey had to compare 2 object stimuli and indicate whether these were the same or different. Each trial contained 2 object stimuli presented successively. Two white squares were presented after the second stimulus at 10-degree eccentricity left and right on the display. The monkey was rewarded for making an eye movement to only 1 of these squares (left if the 2 stimuli were the same, right if the 2 stimuli were different). The monkeys learned the task procedure initially with stimuli presented at the same screen location (center of gaze) with no interstimulus interval (ISI), and gradually, monkeys learned to perform the task with a random position jitter of as much as 2 visual degrees and with an ISI of at least 420 ms (stimulus duration of 300 ms for each object). Before we started scanning, the monkeys were able to perform the task ~90% correct even on the first block of 240 trials with sets of easily discriminable stimuli that they had never seen before. In this task, stimulus presentation was dependent on the monkey's behavior—stimulus presentation was not started unless the monkey was fixating and was aborted if fixation was broken during the trial.

Shape Discrimination Task: Training to Learn Within-Object–Class Discriminations

Scanning was interrupted in each monkey for several months (J: 95 days; M: 81 days). During this period, each monkey was trained intensively to discriminate among exemplars within 1 of the 3 object classes (J: spikies; M: smoothies). Although each monkey was able to perform the "shape discrimination task" at 90% with new sets of easily discriminable objects, they initially performed at chance for discriminating among exemplars within the novel object class (J: 50% correct; M: 53% correct), even though we used the largest possible stimulus differences (maximum difference on each of the shape dimensions within the object class). We tested this initially with ISI 420 ms and a stimulus position jitter of 2 visual degrees. We made the task easier with an ISI of 0–10 ms and no position jitter, and after a few daily sessions, each monkey had above-chance performance (on session 3, J: 63% correct; M: 66% correct). After a few months of training with slowly increasing ISI and position jitter, each monkey attained good performance for the largest stimulus differences with ISI 420 ms and a position jitter of 2 visual degrees (J on session 40: 75% correct; M on session 36: 73% correct). We continued training each monkey for smaller stimulus differences (either no difference on some shape dimensions or smaller-than-maximal differences on several shape dimensions). When scanning started again with the (initially) novel object classes, each monkey attained good performance for stimulus differences of only 60% of the maximal differences (J on session 46: 71% correct; M on session 43: 72% correct). By this time, each monkey had seen more than 130 000 presentations of exemplars from (only) the trained object class in the context of this shape discrimination task. For comparison, the scanning period before this behavioral training period involved less than 6000 presentations of exemplars from each object class. For the fMRI experiments following this behavioral training period, we trained each monkey regularly on days in between scan sessions. We performed all behavioral training in the mock scanner, but we confirmed in the last stage of the training phase that shape discrimination performance after training was very similar in the scanner (tested with 75% of the maximal stimulus differences: performance in the mock scanner and in the scanner was 75% and 73%, respectively, for monkey J, and 81% and, 77% respectively, for monkey M).

Shape Discrimination Task: Scanning

For one experiment (Experiment 3, see further), we adapted the shape discrimination task for use in the scanner by making the stimulus presentation independent of the monkey's behavior. This adaptation also involved a reduction of the intertrial interval to increase the number of stimuli presented per unit of time (to 24 trials and 48 stimuli per 45 s). As shown in the previous section, the monkeys were able to perform this task very well with exemplars of the trained object class. However, prior to Experiment 3, monkeys did not perform the shape discrimination task for exemplars from the 2 other object categories. Not surprisingly, each monkey performed much better ($P < 0.001$, *t*-test across time series) for the trained objects (J: 74% correct; M: 71%) than for the untrained objects (J: 51% correct; M: 50% correct). To counter the possibility of motivational problems as a consequence of the lower reward levels in the blocks with untrained objects, monkeys were also rewarded on some error trials in these blocks (38% of error trials) in blocks with untrained objects in order to keep the overall level of reward the same in the different stimulus conditions. The overall level of reward was still higher on correct trials than on error trials, so the monkeys would still benefit from doing the task as well as possible. To exclude the possibility that monkeys might be attending less during blocks with untrained stimuli, we inserted random "catch trials" in 1 behavioral session of monkey J. This monkey was trained with spikies, and 1 out of 8 trials (never the first) in blocks with smoothies and cubies was a catch trial with spikies. Performance for the spikies was 82% on these catch trials ($n = 96$), similar ($P = 0.22$) to the performance of 76% in blocks with trained objects and significantly higher ($P < 0.0001$) than the performance with untrained objects (51%).

fMRI Scanning and Experimental Manipulations

Scanning was carried out at the Martinos Center for Biomedical Imaging at Massachusetts General Hospital in a 3T Siemens Allegra magnet (Siemens, Erlangen, Germany) with a radial surface coil that was positioned immediately over the head (Tsao et al. 2003). Functional images were acquired with an echo planar imaging (EPI) sequence (135 time points per time series; repetition time [TR], 3 s; echo time [TE], 26 ms; 64 × 64 matrix; voxel size 1.25 mm isotropic; 45 slices). Slices were approximately coronal and covered the entire ventral visual pathway from the occipital pole to the temporal pole. We also performed 1 anatomical scan session in each monkey under general anesthesia (combination of ketamine and xylazine). The anatomical resolution of these anatomical volumes (3-dimensional magnetization

prepared rapid acquisition gradient echo volume) was 0.50 mm isotropic, and we acquired at least 5 volumes in each animal to increase signal-to-noise ratio by averaging.

All fMRI data were collected using a block design in which each block (45 s, see below) contained exemplars from only 1 object class. Each monkey was initially scanned for just a few sessions without contrast agent (blood oxygen level–dependent [BOLD] scan sessions) to solve any remaining technical and behavioral problems. These BOLD scans always involved the localizer stimulus conditions (monkey faces, objects, and Fourier-scrambled images). All subsequent scan sessions were preceded by an i.v. injection of a monocrystalline iron oxide nanoparticle (MION) contrast agent (around 7 mg/kg) (Vanduffel et al. 2001; Leite et al. 2002). For each experiment, we performed multiple scan sessions (i.e., multiple days) that each included multiple time series consisting of 3 stimulus conditions (2 blocks of each stimulus condition per time series) and a fixation condition (3 blocks per time series). Thus, a total of 9 blocks were performed in each time series, and each block lasted for 45 s (15 time points), for a total time of ~7 min in each scan run (time series). Unless noted, visual object stimuli were always presented at the fovea with a horizontal and vertical position jitter of ±1 visual degree (from stimulus to stimulus), for 300-ms duration, and with a 420-ms ISI. In total, 6 experiments were carried out during MION scan sessions.

Experiment 0
Localizer experiment (27 and 31 time series in monkey J and M, respectively). The 3 stimulus conditions in this experiment were monkey faces, objects, and Fourier-scrambled images. Monkeys always performed the color task. Some of these data were collected before Experiment 1 and some were collected after Experiment 1 and after the shape discrimination training phase.

Experiment 1
Novel object experiment "before" shape discrimination training (27 and 26 time series in monkey J and M, respectively). The 3 stimulus conditions in this experiment were smoothies, spikies, and cubies. Sixty fixed exemplars of each class were presented in a random order in a stimulus block of 45 s. These 60 exemplars were chosen to span the total space of exemplars in each class, and they were the same in each block. Monkeys always performed the color task. The primary experimental goal was to uncover any spatial patterns of selectivity for the 3 initially novel object classes within IT. The data from this experiment are referred to as "original" in the figures.

Experiment 2
Novel object experiment "after" a 3-month time interval with shape discrimination training (28 and 33 time series in monkey J and M, respectively). Same procedures as for Experiment 1. This experiment was identical to Experiment 1 except that it followed 3 months of extensive shape discrimination training (see Behavioral Tasks, above). The experimental goal was to test the across-time and across-training stability of the object-class–selective responses found in Experiment 1.

Experiment 3
Shape discrimination task experiment (35 and 26 time series in monkey J and M, respectively). The 3 stimulus conditions in this experiment were smoothies, spikies, and cubies. The monkeys performed the shape discrimination task adapted for use in the scanner (see above in section Shape Discrimination Task: Scanning). The experimental goal was to test the across-task stability of object-class–selective responses.

Experiment 4
Position experiment (43 and 27 time series in monkey J and M, respectively). The 3 stimulus conditions in this experiment were smoothies, spikies, and cubies. In all other experiments, the stimuli were positioned at the fovea with some jitter (see above)—here we further varied the retinotopic position of the stimuli. In monkey J, we centered the position of the stimuli to very eccentric positions (8.4 visual degrees) so that the stimulated visual field area was not overlapping at all with the visual field area stimulated by the foveal stimuli. The eccentric stimuli were restricted to 1 quadrant, and different quadrants were tested in different time series. The monkey performed the passive fixation task (to limit potential eye movements toward the stimuli). In monkey M, we tested the same question by compensating any differences in retinotopic envelope of the different object classes by an object-class–specific jitter in position (e.g., the spikies were jittered much less in the vertical direction than the cubies; see Supplementary Fig. 1). The monkey performed the color task. The experimental goal was to test the stability of object-class–selective responses to changes in stimulus position.

Experiment 5
Global versus local shape features experiment (33 and 37 time series in monkey J and M, respectively). The 3 stimulus conditions in this experiment were spiky smoothies, cuby spikies, and smoothy cubies (see Fig. 9*a*). Monkeys always performed the color task. The experimental goal was to test effects of changes in stimulus shape on the object-class–selective responses.

Importantly, some of the experimental data were acquired in a specific order: Experiment 1 was performed before the 3-month shape discrimination training phase, Experiments 2–5 were acquired after this training phase. All data of Experiment 2 were collected immediately after the training phase; Experiments 3–5 were tested after Experiment 2. We have taken the data of Experiment 1 as the baseline, pretrained, original data set, and the data from all other experiments are compared with this baseline.

Fixation Condition
All experiments included a fixation condition that was used as a baseline. The characteristics of this fixation condition were slightly different in the different experiments in order to match the stimulus conditions as closely as possible (i.e., other than the objects shown). During passive fixation, the display in the fixation condition only showed the fixation target (a small white square) at the center of the screen. During the color task (used in most experiments), the display in the fixation condition showed this fixation target and, in addition, a response target at the top of the screen. Finally, during the shape discrimination (Experiment 3), the fixation target disappeared for regular intervals during which 2 response targets were presented left and right. The timing and the frequency of these events were matched to that used in the shape discrimination condition. In addition, as in the object-class conditions, the monkey was rewarded for making a saccade to 1 of the 2 response targets in the fixation condition (in the fixation condition, the rewarded target was determined randomly for each trial). In all experiments, reward schedules were arranged so that the total amount of reward was similar across conditions.

Statistical Preprocessing
Data were analyzed with Freesurfer (http://surfer.nmr.mgh.harvard.edu/) (Dale et al. 1999; Fischl et al. 1999), froi (http://froi.sourceforge.net), as well as custom Matlab code. All volumes of each time series were aligned to a reference functional volume from one particular scan session (the reference volume was from Experiment 2 in monkey J and from Experiment 1 in monkey M) (Cox and Jesmanowicz 1999). This step combines motion correction and between-session alignment. The reference functional volume was coregistered with the high-resolution anatomy.

Inclusion of Data for Analysis
A time series was retained for analysis if the monkey did not move too much (in general, time series with movement spikes of over 1 mm of total translation did not provide useful data) and if task performance was good (overall percentage fixation in the fixation task and the color task >84%). Even though the animals were head restrained via head post (see Animals and Surgery, above) during scanning, and were trained to remain still, moderate motion of the body and very small motion of the head was still possible. Because the amount of motion and task performance are both likely to depend on animal motivation level, it is not surprising that time series with poor task performance typically showed extensive motion. We retained 193 and 180 time series from 21 and 19 sessions in monkey J and M, respectively (~87% of the time series obtained in these sessions). In each monkey, data

from at least 26 time series were included in each experiment (acquired in 3–4 scan sessions per experiment per monkey).

Analysis of Eye Movements

Following previously described methods (Vanduffel et al. 2002), we added 3 measures of the monkeys' eye movements to the analyses as covariates (in addition to the 6 motion correction parameters, 3 translations and 3 rotations). These 3 measures were the x and y coordinates of the foveated visual field position and the total deviation of the center of gaze relative to the center of the fixation point. Each of these measures was subsampled to 1 TR (3 s) and convolved with the hemodynamic response. In Supplementary Material, we describe additional analyses to rule out effects in our results due to differences between stimulus conditions in either the monkeys' eye movement patterns or other behavioral indices.

Smoothing, Normalization, and Gamma Fitting

fMRI data were spatially smoothed with a Gaussian kernel of 2.5-mm full width at half of maximum. Preliminary analyses suggested that this degree of smoothing was needed to obtain sufficient signal-to-noise to detect significantly activated voxels and to look at the spatial distribution of activity. Thus, our methods allow us to look at organization with a grain of 2–3 mm. In Supplementary Material, we show that this amount of smoothing maximizes the reliability of the spatial distribution of activity. In addition, the data were scaled to have the same global signal intensity for each time series. The predictor for each stimulus condition (0 or 1 at each time point) was convolved with a gamma function, and the general linear model was used to compute the response of each voxel in each condition. The gamma function was given by the formula:

$$f(t) = (g \times e^2) \Big/ 4 \times ((t-d)/g)^2 \times e^{-((t-d)/g)} \text{ with } g = 8\text{s and } d = 0\text{s}.$$

These parameters were chosen to approximate the impulse response function for MION signal (Vanduffel et al. 2001; Leite et al. 2002). The response for each condition in each voxel is expressed in units of percent signal change (PSC), which is the response in each condition minus the response in the fixation condition (see above), normalized by the mean signal value at each voxel. We reversed the responses in all analyses because with MION higher blood volume is associated with lower signal (Vanduffel et al. 2001).

Selection of Regions of Interest

Significance maps of the brain were computed by performing t-tests for pairwise comparisons of conditions and thresholded at $P = 0.0001$ (uncorrected for multiple comparisons). We used this same statistical threshold throughout all analyses to define the following regions of interest (ROIs):

(i) Novel-object–selective ROI: All significantly activated voxels in IT cortex from all pairwise comparisons of novel object classes in Experiment 1. All these (often discontinuous) voxels were taken together as 1 novel-object–selective ROI with a corresponding list that indicated the preferred and unpreferred object class for each voxel. If a voxel showed significant selectivity in more than 1 contrast, then it appeared more than 1 time in this list (with a different preferred or unpreferred class; so as different voxel by class combination). A novel-object–selective ROI was created for each scan session in Experiment 1 (average of 291 voxel by class combinations per session and average of 223 individual voxels per session). For each time series, we calculated the response to the preferred and to the unpreferred object class averaged across all voxels in each of the 3 session-specific ROIs. For one of the sessions, the time series was part of the data used to select the ROI, so this may lead to an overestimation of the true selectivity in this voxel. The responses in a time series averaged across the 2 other sessions provide an independent estimate of how well object-class selectivity is replicated across scan sessions (see Fig. 4). The amount of selectivity reported in the main text corresponds to this independent estimate of selectivity, and it is also the selectivity measure that is compared with the selectivity found in the other experiments (see further section Statistical Analyses I: Comparing Selectivity in the Voxels with Strongest Selectivity).

(ii) Standard-object–selective ROI: All voxels in IT cortex with significant activation in the contrast (objects—scrambled) in Experiment 0 (average of 1264 voxels per monkey). The definition of this ROI is comparable to previous work in monkeys and humans (Malach et al. 1995; Tsao et al. 2003; Denys et al. 2004). This ROI includes voxels that are activated by a variety of familiar objects more than by unstructured scrambled images and is thus unbiased with respect to selectivity among the 3 novel object classes. This ROI was used to calculate the correlations in the spatial distribution of activity across IT cortex evoked by novel objects.

(iii) Visually responsive ROI: All voxels ventral to the dorsal bank of the superior temporal sulcus (STS) along the ventral visual pathway (occipital pole to temporal pole) with significant activation in the contrast of all stimulus categories (faces, objects, and scrambled images) versus the fixation baseline in Experiment 0 (average of 8271 voxels per monkey). This large ROI was mainly used for display purposes, more specifically, to select the voxels to which the combined color scale (see further) would be applied in Figures 5, 6, 8, and 10.

(iv) Parafoveal V1 ROI: All voxels from the visually responsive ROI that were located around the anatomical location of the parafoveal visual field (up to ~7 visual degrees) in V1 (average of 1076 voxels per monkey). We did not perform retinotopic mapping in these 2 monkeys because our study deals with high-level visual cortex and because defining part of V1 based on anatomical landmarks is straightforward given the consistent anatomical location of this area as found in previous studies (Brewer et al. 2002; Fize et al. 2003). This ROI was used to compare the sensitivity to stimulus position manipulations in early visual cortex with that observed in IT cortex. Similar results were found with more inclusive ROIs of early visual cortex, for example, all voxels in the visually responsive ROI with an anterior/posterior position more than 7 mm posterior to the external auditory meatus.

The novel-object–selective ROI and standard-object–selective ROI were both restricted to IT cortex so that any significantly activated voxels outside this region were not selected. IT cortex was defined as the temporal lobe region lateral and inferior to the fundus of the STS (including the lower but not the upper bank of the STS). The most posterior point was ~2 mm anterior to the external auditory meatus (including area TE but not area occipitotemporal area [TEO]), and the most anterior point was ~22 mm anterior. The medial boundary on the ventral surface was ~10 mm lateral of the midline. The estimated border between area TEO, an area with a crude retinotopic organization (Boussaoud et al. 1991), and area TE, an area without any known retinotopic organization, is shown in Figure 3. The as such defined IT region encompassed about 2600 functional voxels across both hemispheres. Almost all these (average of 2569 voxels) were also part of the visually responsive ROI. Almost all the voxels in the novel-object–selective ROI (94%) and standard-object–selective ROI (94%) were included in the visually responsive ROI. There was less overlap between the novel-object–selective ROI and the standard-object–selective ROI: of the voxels included in the novel-object–selective ROI, 54% (monkey J) and 83% (monkey M) were included in the standard-object–selective ROI.

Statistical Analyses I: Comparing Selectivity in the Voxels with Strongest Selectivity

We defined the novel-object–selective ROI as the population of voxels that showed significant selectivity for novel object classes in a single scan session of Experiment 1. For these voxels, we quantified selectivity for the novel object classes using the independent data from the other sessions of Experiment 1. For each time series, we calculated the response (PSC) to the preferred object class and to the nonpreferred object class for each voxel; we then averaged the responses across voxels (each of which can have a different preferred and nonpreferred object class). This gives 2 values for each time series, the response to the preferred category and the response to the nonpreferred category in that time series. To compare selectivity across experiments, we first divided these data of each time series by the average response (averaged across all voxels) to the preferred category in the experiment to which the time series belonged. Thus, the response to

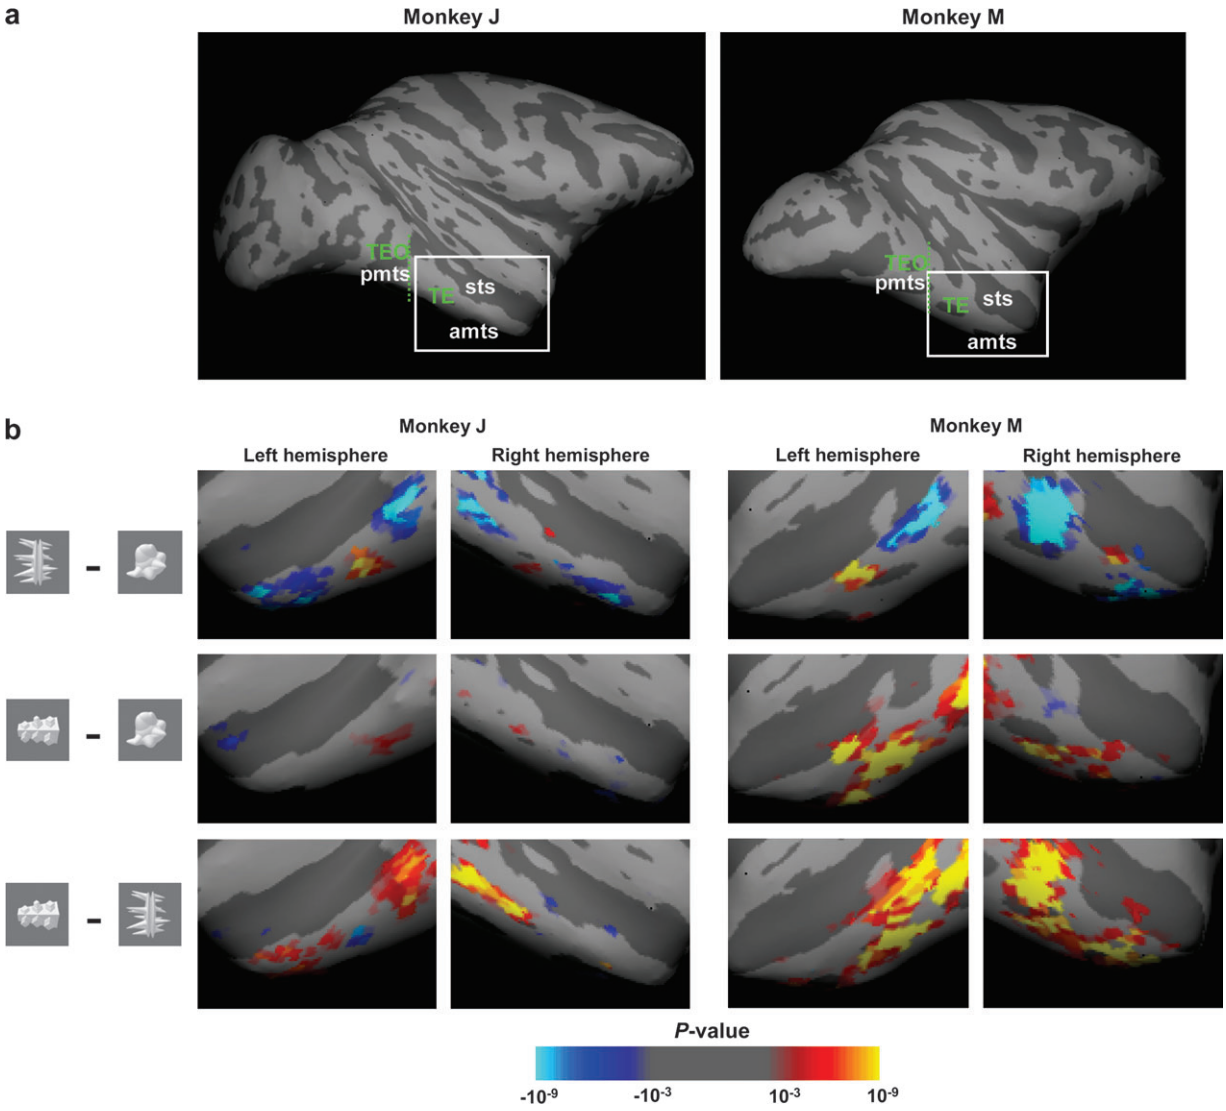

Figure 3. Significance maps of selectivity for novel object classes across anterior IT cortex (data from Experiment 1). (*a*) Anatomical location of anterior IT cortex (area TE). The border between area TEO, an area with a coarse retinotopic organization (Boussaoud et al. 1991), and area TE, an area without any known retinotopic organization, is indicated with a green dotted line and was estimated based on data in the literature (Boussaoud et al. 1991). Abbreviations of anatomical landmarks: pmts, posterior middle temporal sulcus; amts, anterior middle temporal sulcus. (*b*) Significance maps in IT cortex for each of 3 pairwise contrasts ([spikies — smoothies], [cubies — smoothies], [cubies — spikies]) in each hemisphere of each monkey. The maps are thresholded at $P < 0.001$ (uncorrected for multiple comparisons).

the preferred category was normalized to 1.0 for each experiment (this is the "normalized fMRI response" shown in Fig. 4). This normalization factor was very reliable and significantly above 0 in each case where it was used. We expressed selectivity as: (preferred – nonpreferred) × 100%. Due to the normalization, this "selectivity index" reflects the strength of the response to the nonpreferred object class relative to the response to the preferred object class. The normalization insures that the selectivity index reflects true selectivity, independent from overall responsiveness. The latter might be affected by various technical details, such as the batch/concentration of MION used and the exact distance between the surface coil and IT cortex. Note that this normalization was only done for the ROI analyses (Figs 4 and 7) and not for other analyses.

We used a 2-tailed *t*-test to determine whether the selectivity was significantly different from 0 given the variability of the selectivity across time series. A significant selectivity means that the preference of voxels selected for a particular object-class pairwise preference (based on data from 1 scan session of Experiment 1, see above) was replicated in independent data obtained from other scan sessions. In addition, we used an unpaired 2-tailed *t*-test to determine whether the selectivity in any of the other experiments was significantly different from the selectivity in Experiment 1. In general, we performed these statistical tests for each monkey separately. When we refer to *t*-tests that combine the data of the 2 monkeys, then this was done by considering all time series together ("fixed-effect" analyses).

Statistical Analyses II: Correlational Analysis of Spatial Distribution of Activity

We used correlational analyses to determine the stability of the spatial distribution of category-selective activity across experiments. For each voxel in the ROI and each experiment, we calculated 3 pairwise selectivity indices: [smoothies – spikies], [smoothies – cubies], and [spikies – cubies]. For each of these indices, we calculated the across-voxel correlation between the original condition (Experiment 1) and the other experiments. The higher these correlations, the better the correspondence between experiments in the spatial distribution of object preferences. The correlation values presented in the main text refer to these across-voxel correlations when all relevant data from each experiment are first averaged together (within experiment)

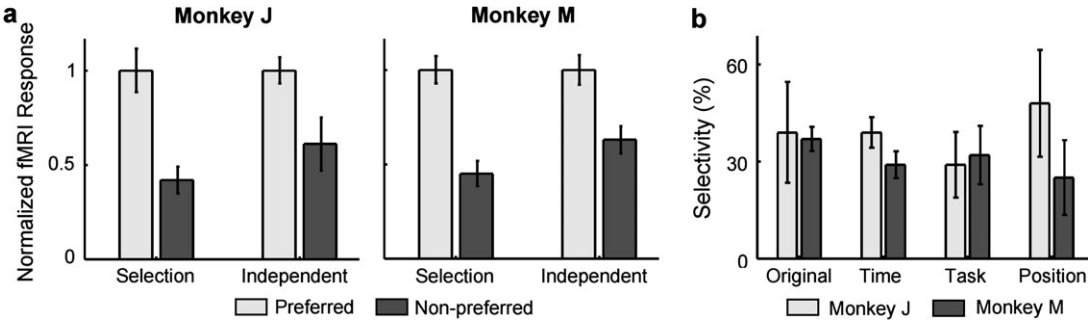

Figure 4. Reproducible, stable selectivity in monkey IT cortex for novel object classes. (*a*) The mean response to the preferred and nonpreferred object classes over voxels showing significant pairwise object class selectivity in Experiment 1(normalized to the mean response to the preferred object class, see Materials and Methods). The difference between the light and dark bar indicates the average pairwise class selectivity. Responses are shown for the scan session used to select the voxels with a significant preference for one object class above another (selection) and for the same voxels but using data from other scan sessions (independent). (*b*) The category selectivity in these voxels, defined as normalized preferred response minus normalized nonpreferred response. This selectivity measure was computed for the selected voxels using independent data from Experiment 1 (i.e., the independent data from *a*; original), data obtained several months later (time), data obtained while the monkeys performed a different visual task (task), and data obtained with objects presented at different visual field positions (position). Error bars show the standard error of the mean across time series (see Materials and Methods).

before computing the spatial correlations across experiments. For each comparison (e.g., Experiment 1 with Experiment 2), we assessed the variability of these correlations that is due to intrinsic variability in the data by computing the correlation of the averaged data of Experiment 1 with the data from each individual time series of the other experiment (e.g., Experiment 2). We determined whether these correlations were significantly higher than 0 with a 1-tailed *t*-test (degrees of freedom = number of time series – 1). The error bars in Figures 6*b* and 9*b* show the standard error of this correlation obtained with individual time series, multiplied by the ratio of the correlation when all time series are averaged to the mean correlation obtained with the individual time series (of course, correlations are lower for single time series than after averaging). Thus, the correlation values shown in Figures 6*b* and 9*b* divided by the shown error bars reveal the *t* value (the test statistic).

These statistical tests assess whether correlations are significantly positive. A related question is how much of the explainable variance is captured by a positive correlation. Because the data are intrinsically noisy, only an infinite amount of data (averaged in each condition) would produce a correlation value of 1.0, even when an experimental manipulation has no effect on the spatial pattern of selectivity. In a statistical sense, some fraction of the variation in the observed selectivity topography is due to reproducible differences in voxel preferences across the object classes (i.e., signal), and the remaining fraction is due to noise intrinsic to the methods (i.e., noise—variation that does not reproduce from time series to time series). The former fraction of the variation—the, so-called, explainable variance—provides a measure of the maximum correlation that one can expect to see when comparing across 2 experiments (i.e., the maximum observed correlation to be expected if there is no underlying difference in the selectivity topography evoked in the 2 experiments). The formulae to compute the explainable variance have been developed in basic measurement theory (Cronbach 1949; Lewis-Beck 1994) and are explained in more detail in Supplementary Material. In short, we divided the data of each experiment in half, and we computed the spatial correlation between those 2 halves for each pairwise selectivity index. This correlation provides a measure of the maximum correlation that can be expected (after correcting for the fact that it is based on only half the data), thus giving an estimate of the reliability of the full data set. The reliability of each of the 2 full data sets (e.g., Experiment 1 and Experiment 2) can then be used to determine the maximal expected correlation between the spatial distribution of selectivity observed in each experiment. The statistical test of the difference between observed correlations (across Experiments) and the maximal expected correlation was again based on the variability of correlations across time series of each experiment by correlating the averaged data of Experiment 1 with the data of individual time series from the other experiments. We determined whether the single time series correlations were significantly lower than the reliability of the involved data sets with a 1-tailed *t*-test (degrees of freedom = number of time series – 1).

Color Maps of the Spatial Distribution of Activity

To display the spatial distribution of activity on inflated brain images, we used Freesurfer functional analysis stream to compute 3 pairwise *t* values for all visually responsive voxels ventral to the STS. A *t* value is a measure of selectivity that takes into account both the difference in response between 2 conditions as well as the variability of this difference across time points. We used standard Freesurfer commands (paint-sess and surf-sess) to display this *t*-map on a lateral view of the inflated brain for each pairwise contrast with the standard blue to red/yellow 1-dimensional color scale. In Figure 3, significance maps are shown for each pairwise comparison with a threshold of $P < 0.001$, uncorrected for multiple comparisons. In later figures that illustrate the correlation analyses, we used a combined color scale to plot selectivity for all pairwise comparisons at once. These color maps (Figs 5, 6, 8) are not intended to show individually significant voxels (e.g., as in Fig. 3) but to illustrate the pattern of object-class preference across IT cortex. For the pairwise *t*-maps that were used as the basis for this combined map (illustrated in Fig. 5*a*), the parameters for the standard blue to red/yellow color scale were selected so that the threshold was 0 and so that the colors related to maximal selectivity corresponded to *t* values –5 (blue) and 5 (yellow). All absolute *t* values higher than 5 are shown with the same "maximal selectivity" colors. The *t*-map images obtained for the 3 pairwise *t*-maps were then combined into 1 image with a multidimensional color scale so that hue represents which object class is preferred and saturation the amount of selectivity as represented by the *t* values (no saturation = gray = *t* value near 0 in each contrast). To do this, the preference for each object class was associated with the value of 1 of the RGB color channels, and the values ranged from 128 to 256. For example, preference for smoothies was associated with the red channel. If a voxel showed no preference for smoothies in any contrast involving smoothies, then the value of the R channel was set to 128. If a voxel showed a preference for smoothies with a *t* value equal or greater than 5 in 1 contrast and not in the other contrast involving smoothies, then the value of the R channel was set to 192 (=128 + 64 × 5/5). If a voxel showed a preference for smoothies with a *t* value equal or greater than 5 in each contrast involving smoothies, then the value of the R channel was set to it maximum value of 256 (=128 + 2 × 64 × 5/5). If a voxel showed a preference for smoothies with a *t* value equal to 1.5 in each contrast involving smoothies, then the value of the R channel was set to 166 (=128 + 2 × 64 × 1.5/5). If all 3 channels were set to 128, the resultant color would be gray—indicating no preference among the 3 object classes.

The significance of a *t* value is related to the number of time points per voxel, which was different for different experiments. In the experiment with the lowest number of time series, a *t* value of 5 corresponded to an uncorrected *P* value of $P < 10^{-6}$ (*t* value of 5 in 1 of the contrasts results in a color from the outer border in the color scale of Fig. 5). A *t* value of 2.5 (color is average of gray and maximal saturation) corresponds to an uncorrected *P* value of $P < 0.02$. Supplementary Figure 3 shows color maps thresholded at $P < 10^{-3}$ (uncorrected).

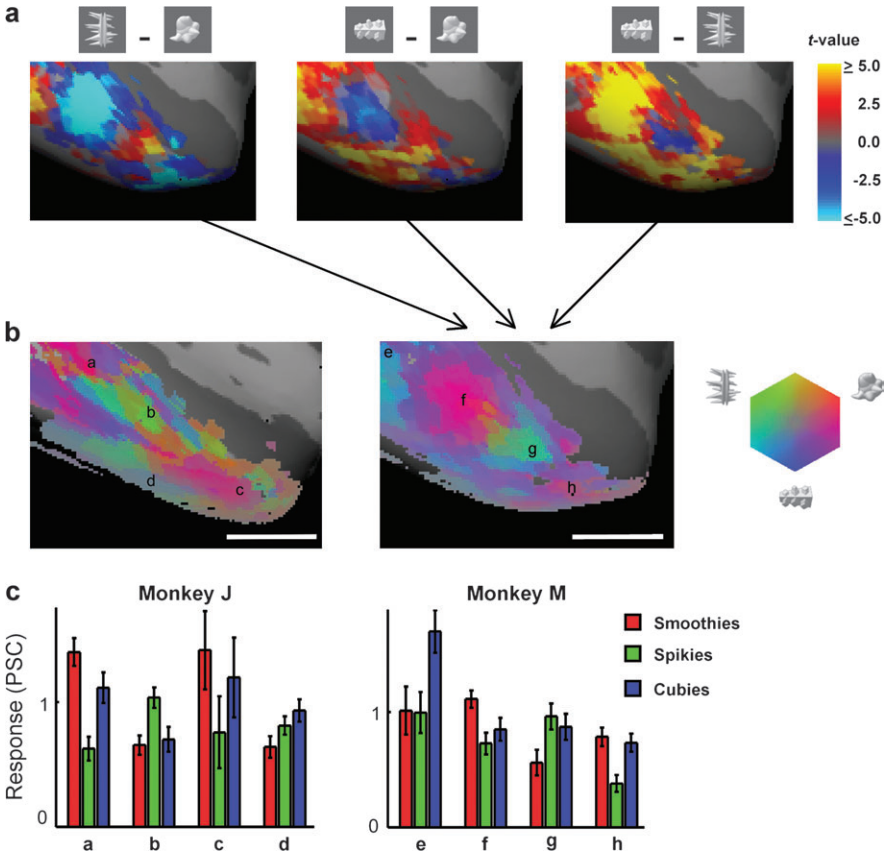

Figure 5. Large-scale selectivity maps for novel objects across the ventral visual pathway in monkeys. (*a*) Nonthresholded pairwise *t*-maps of the 3 object classes in monkey M (right hemisphere). (*b*) Multidimensional color map that represents the selectivity across all 3 object classes (right hemisphere of each monkey). The procedure followed to derive these color maps from pairwise *t*-maps (as shown in panel *a*) is explained in the Materials and Methods. Saturation represents amount of selectivity (no saturation = gray = same response to each object class) and hue represents which object class is preferred (spikies = green, smoothies = red, cubies = blue) or which combination. Scale bar shows an approximate distance of 7 mm. (*c*) Mean response to each object class at points indicated in the selectivity maps in *b*. Responses are expressed in units of PSC relative to the fixation condition, and error bars represent the standard error of the mean of these responses across time series (see Materials and Methods).

Results

Here we present results regarding the spatial pattern of the fMRI selectivity for 3 initially novel object classes (smoothies, spikies, and cubies, see Fig. 2) across IT cortex. First, we used data collected from the very first day the subjects saw the object classes and standard methods to reveal that IT cortex contains a subset of voxels that "individually" show reproducible, highly significant selectivity among the 3 novel object classes. However, we found that these voxels are just the tip of the iceberg, and we proceed to show that, using correlation methods on the same data, IT cortex contains distributed and reproducible spatial patterns of selectivity. Second, we performed a series of experimental manipulations to test the stability of this selectivity pattern across time and training, across behavioral task, and across changes in object position. Third, we tested manipulations of stimulus shape by changing the global and local shape properties of the object classes. Finally, we describe the relationship between this selectivity pattern and the previously described face patches in IT cortex. All experiments were designed, a priori, to target IT cortex, while using the earlier visual areas as controls. Nevertheless, as described in Supplementary Material, we found no stable patterns of selectivity for these novel object classes in other brain regions that have previously been shown to prefer intact object images over scrambled images.

Reproducible Topography of Selectivity for Novel Object Classes across IT Cortex

In each of 2 monkeys that were naive to the object classes before the first fMRI scans, fMRI scanning revealed a reliable spatial distribution of object-class selectivity spanning IT cortex (Experiment 1). That is, different IT voxels preferred different object classes, a finding that we refer to as a "topography of selectivity" throughout this text.

We began our analyses by using standard methods to determine if any IT voxels were individually statistically selective among the 3 novel object classes. Figure 3 shows the significance maps across IT in each pairwise comparison of novel object classes. These maps illustrate that some voxels in IT are selective when considered individually (an average of 19% of IT voxels per pairwise comparison when thresholded at $P < 0.001$, uncorrected). We obtained an unbiased measure of the replicability and strength of this object-class selectivity across scan sessions (separated by 8 days on average) by performing test–retest analyses for voxels with strong object-class selectivity (see Materials and Methods). In particular, we selected IT voxels with a significant preference for one object class over another in 1 scan session (novel-object–selective ROI described in the Materials and Methods). We found that these same voxels showed a similar preference in other, non-overlapping scan sessions (Fig. 4*a*), and this between-session

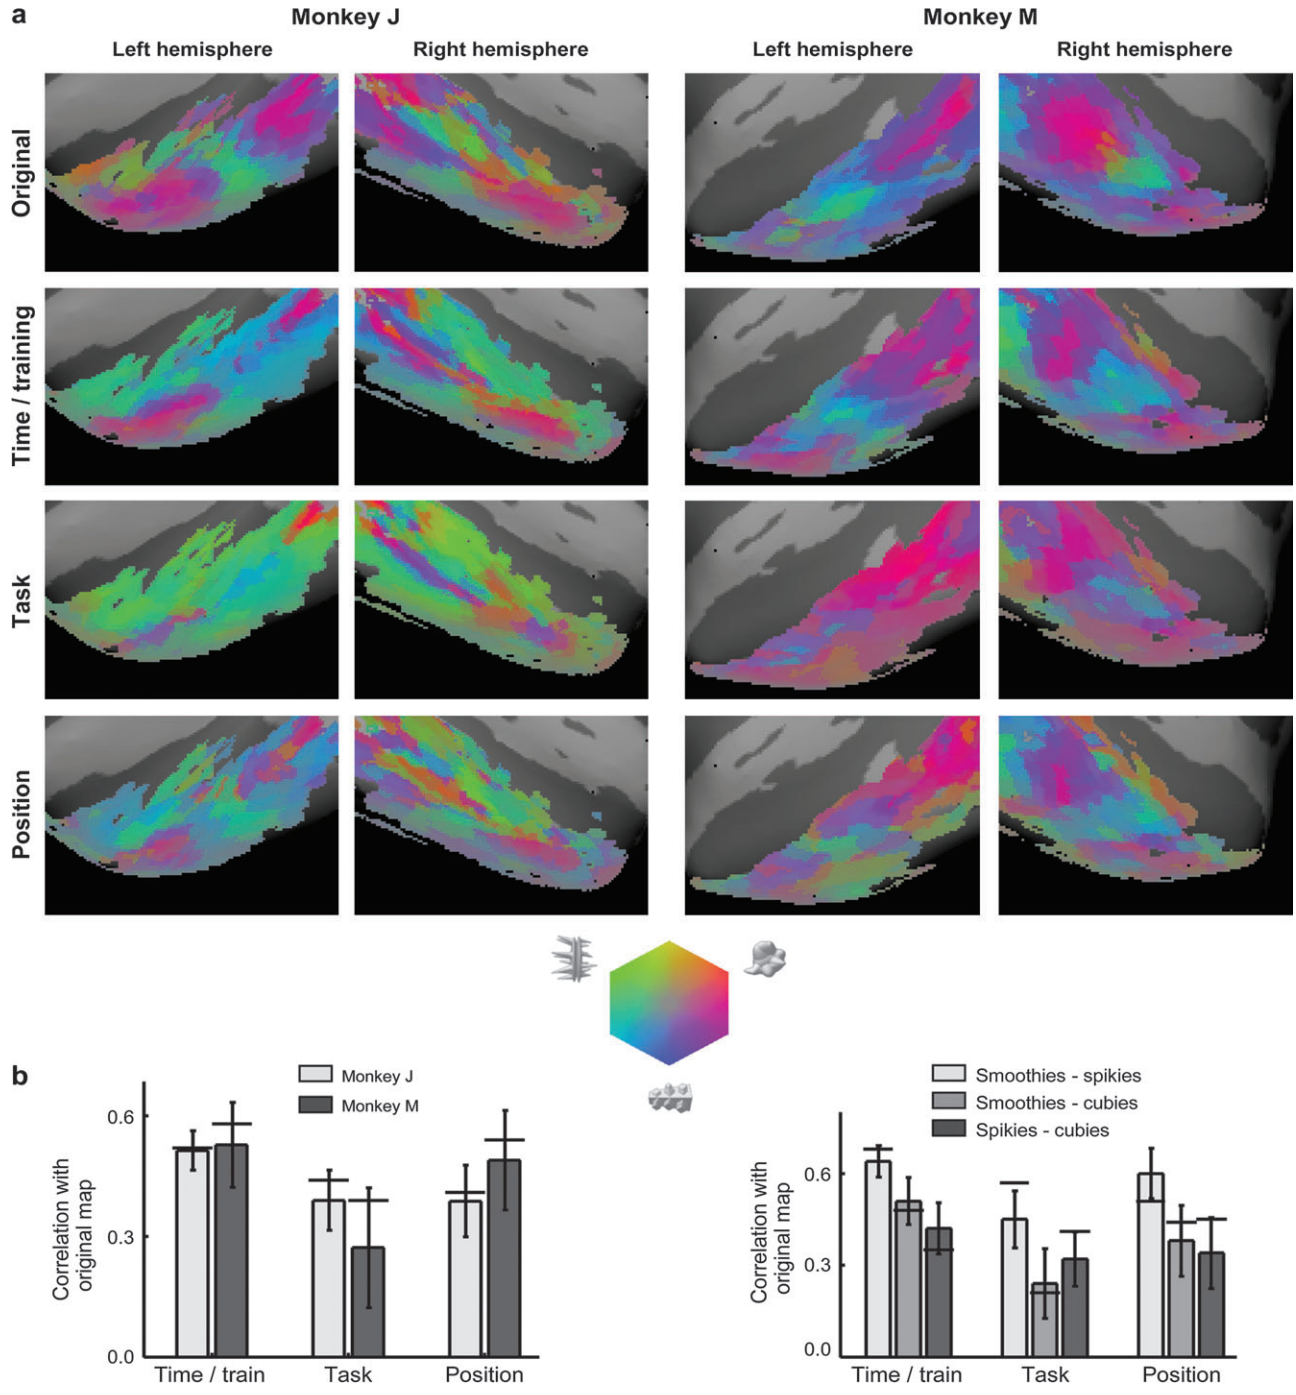

Figure 6. Stability of the selectivity topography in IT cortex across time, training, task, and object retinal position. (*a*) Selectivity topography represented by the same continuous color scale introduced in Figure 5. The maps are shown for the original experiment and across time/training, task, and changes in the position of the stimuli. (*b*) The correlation between the original selectivity topography and the selectivity topography measured from each of the other data sets in each monkey (including data from both the left and right hemispheres; see Results for details). The left panel shows the correlations for each monkey separately, the right panel shows the correlations for each pairwise object class comparison. The horizontal lines represent the maximal correlation that is expected if the selectivity topography was identical in each case (i.e., given the reliability of the 2 data sets being correlated; see Materials and Methods). The ratio of the observed correlation to the error bar represents the test statistic to determine the significance of the correlation across time series (see Materials and Methods).

replication was significant across independent time series in each monkey (J: $P = 0.02$; M: $P < 0.001$, see Materials and Methods). The response of each voxel to its nonpreferred object class was, on average, only 62% of its response to its preferred object class. Throughout this paper, we quantify this selectivity by the object-class selectivity index—the difference in the response to the preferred and nonpreferred object classes times 100 (after normalizing the response to the preferred object class to 1). In Figure 4*a*, the mean object-class selectivity index was 38% (the expected value if no selectivity existed is 0%).

In sum, we found reproducible selectivity in IT among initially novel object classes, at least when we focus on the voxels with the strongest selectivity (above). However, these

regions might only be the "tip of the iceberg." IT cortex might contain a reproducible pattern of selectivity that spans a larger region of IT than that revealed when voxels are only considered individually (Haxby et al. 2001). To investigate this, we used correlation methods to analyze the pattern of selectivity across all voxels in IT cortex with a preference for objects over scrambled images (the standard-object–selective ROI; average of 1264 voxels per monkey; see Materials and Methods). Specifically, in each monkey, we first computed 3 pairwise response difference maps for these voxels (i.e., 1 map for each comparison: [smoothies – spikies], [smoothies – cubies], and [cubies – spikies]) and then computed the reliability of each pairwise map by computing the correlation of these maps across nonoverlapping halves of the data from Experiment 1 (see Materials and Methods). The reliability was 0.57 (standard deviation [SD]: 0.12) in monkey J and 0.73 (SD: 0.08) in monkey M (SD computed across different divisions of the data in halves). This significant reliability ($P < 0.01$ in each monkey) shows that IT cortex contains a highly reproducible pattern of selectivity among the novel object classes. This reliable pattern of selectivity did not depend solely on the aforementioned voxels with significant selectivity when considered individually (the voxels used for the test–retest analyses in Fig. 4). In particular, when we removed these voxels from the correlation analysis, then the reliability of the selectivity map in the nonselective voxels was nearly as high (0.55, SD: 0.13 in monkey J, and 0.69, SD: 0.11 in monkey M; $P < 0.01$ in each monkey). Thus, even across all these "nonselective" voxels, we found a statistically reliable pattern of selectivity at the voxel "population" level.

Because our correlation analysis revealed that the fMRI data contain more selectivity than was uncovered in Figure 3, we used a color map method applied to unthresholded maps to illustrate the topography of selectivity across IT cortex. Each potential pattern of class selectivity is captured with a unique color, and Figure 5 illustrates how the color map was constructed from pairwise *t*-maps (also see Materials and Methods). Given that these maps are not thresholded, the goal of these color maps is not to show the significance of individual voxels but to visually illustrate the pattern of object-class preferences across IT cortex and its reproducibility across experiments (see further); in this sense, they resemble the color maps generated with optical imaging techniques. Thresholded maps that only reveal the "tip of the iceberg" of voxelwise reproducible selectivity are shown for each experiment in Supplementary Figure 3.

Figure 5*b,c* shows the color maps for 1 hemisphere in each monkey together with the response to each object class for a few locations in the color map. These responses illustrate that the color maps represent relative preferences for specific object classes rather than a complete segregation of the voxels responding to different object classes. This observation is consistent with the test–retest analyses (see above) that showed that even the most selective voxels did not respond only to a single object class—voxels that preferred one object class also responded to other object classes (relative to a no-object fixation condition), albeit less strongly.

The Topography of Selectivity Is Stable across Time and Training

We scanned both monkeys again after a 3-month interval, using the same methods and behavioral task. During the 3-month interval, each monkey received extensive behavioral training in discriminating among exemplars within just 1 of the 3 object classes (monkey J trained with only spikies, monkey M trained with only smoothies; see Materials and Methods). The other 2 object classes were not seen during this interval. The goal of this training was to do our best to impart special meaning to just 1 of the object classes and to improve the monkey's expertise within that class. We did not aim to distinguish among these and other potential effects of this behavioral training but first aimed to see if such training might alter each monkey's IT pattern of object-class selectivity described above.

Stability across Time

We found that our fMRI measurements of object-class selectivity were remarkably stable across the 3-month training interval. The same voxels selected above for the test–retest analyses (selected using data before the training interval) had a mean selectivity index for the same object class that was nearly identical after the 3-month interval (34%) as computed before the interval (38%, see above; Fig. 4*b*). This selectivity index was highly significant in each monkey (J: $P < 0.001$; M: $P < 0.001$, *t*-test across time series), and in neither of the monkeys was it significantly different from the originally obtained selectivity index (J: $P > 0.5$; M: $P > 0.15$, *t*-test across time series). These numbers are averaged across all 3 object classes, trained and untrained, but similar results were found if the analyses (both for the selection of voxels with significant selectivity and for the computation of selectivity) were restricted to pairwise stimulus comparisons that involved the trained object class either as a preferred or a nonpreferred class (object-class selectivity index of 35% before training and 33% after training).

Although this analysis of selectivity indicates that the "macro" magnitude of selectivity for the novel object classes in the most selective voxels is roughly the same following the 3-month training interval, it does not provide direct insight into the stability of the spatial pattern of selectivity across IT cortex. To examine this, we plotted the selectivity topography that we measured following the 3-month interval (using the same color scheme as in Fig. 5), and we found that the color maps were strikingly similar to those obtained 3 months earlier (Fig. 6). To quantitatively compare the spatial pattern of selectivity across the training interval, we computed the spatial correlation of the selectivity topography maps. Specifically, in each monkey, we again computed 3 pairwise response difference maps for all voxels in the standard-object–selective ROI and then computed the correlation of each pairwise map across the 3-month interval. The correlation (averaged across the 3 pairwise maps) was 0.51 in monkey J and 0.53 in monkey M (see Fig. 6*b*), and it was strongly significant in each monkey (J: $P < 0.001$; M: $P < 0.001$, *t*-test across time series, see Materials and Methods). The correlation across the training interval was also significant for each object pair map in each monkey (monkey J—[smoothies – spikies]: $r = 0.72$, $P < 0.0001$; [smoothies – cubies]: $r = 0.44$, $P < 0.0001$; [cubies – spikies]: $r = 0.38$, $P < 0.0001$; monkey M—[smoothies – spikies]: $r = 0.56$, $P < 0.001$; [smoothies – cubies]: $r = 0.57$, $P < 0.01$; [cubies – spikies]: $r = 0.45$, $P < 0.001$). Thus, we found a reliable pattern of selectivity for each of the 3 possible pairwise contrast maps. Furthermore, the average correlation (0.51) was not significantly smaller than the maximum correlation of 0.55 that is expected if the spatial distribution of selectivity was not changed at all during the interval of 3 months (J: $P > 0.4$; M: $P > 0.4$; the expected correlation takes into account the variability in each data set,

see Materials and Methods and bar plots in Fig. 6*b*). In sum, within the limits of our data, we found that the spatial distribution of activity across IT cortex was remarkably stable across the 3-month training interval.

A drawback of both the analysis of the macro magnitude of selectivity (selectivity index) and the spatial correlation analysis described above is that multiple scan sessions (days) were needed to provide reliable results (the analyses rely on across-session replication of selectivity). Thus, such analyses do not necessarily rule out changes in selectivity that might have happened very rapidly (e.g., over a period of a few days), and for most of our results, the denotation of the 3 object classes as "novel" should thus be taken as a relative concept with this time scale in mind. Furthermore, only the performance of monkey M was good enough in the first scan session with the novel objects to include the data in the analyses (monkey J did not fixate well enough in this first session; see Supplementary Material). Thus, it is possible that these analyses missed relatively rapid changes in the topography of selectivity for novel objects. However, more specific analyses suggest that is unlikely. In particular, the selectivity topography computed from only the first scan session with these novel objects in each monkey (i.e., in monkey J using data with relatively poor fixation performance—around 80%) was already highly correlated to that obtained after the 3-month interval ($r = 0.58$ on average; cf., $r = 0.52$ for all data from Experiment 1, see above). This suggests that the selectivity topography shown in Figures 5 and 6 was already present on the first day that the monkeys saw the novel objects.

Stability across Training

To further investigate possible effects of shape discrimination training, we also focused on the spatial distribution of activity for the trained object class, independently from the spatial distribution of activity for nontrained objects. We used the responses to the object-class cubies as a baseline because this class was not trained in either monkey. This provided us with 2 measures of selectivity per voxel, 1 for trained objects (the response to the trained class minus the response to the cubies class) and 1 for nontrained objects (the response to the other class minus the response to the cubies class; for monkey J: smoothies [nontrained] – cubies [nontrained]; for monkey M: spikies [nontrained] – cubies [nontrained]). With this procedure, the trained and nontrained object classes were perfectly counterbalanced across monkeys. There was no consistent difference in the replicability of the spatial distribution of selectivity for the trained class compared with the nontrained class—monkey J: trained $r = 0.38$ and nontrained $r = 0.44$; monkey M: trained $r = 0.57$ and nontrained $r = 0.45$. Thus, the spatial distribution of activity associated with initially novel objects was as stable across months for an object category that was not seen during this time interval as it was for an object class that was shown more than 130 000 times in a task in which the within-class shape differences were learned.

This far, all analyses focused on the pattern of responses across IT cortex but have ignored potential differences in the "grand mean" of the response to trained and nontrained objects. We tested for training-related changes in overall response in the standard-object–selective ROI with a 2-tailed unpaired *t*-test across time series that compared the posttrained difference in response between trained and nontrained objects with the pretrained difference in response between trained and nontrained objects (to control for possible preexisting differences in responsiveness not due to training). This analysis revealed no significant effect of training on the response to the trained object class in IT cortex (standard-object–selective ROI)—J: $P = 0.28$; M: $P = 0.27$; *t*-test taking the data of the 2 monkeys together: $P = 0.10$ (Fig. 7). Nevertheless, there was a small trend toward stronger responses for trained than for nontrained objects that was also seen in other brain regions (primary visual cortex and prefrontal cortex; see Supplementary Fig. 4). These analyses focus on relative differences between trained and nontrained objects (normalizing to the nontrained responses), but the absolute, unnormalized responses were very similar before and after training (see caption of Fig. 7). Furthermore, measurements of sensitivity for one object class versus another, normalized to the SD rather than to the mean response, also revealed no differential effect of training on the sensitivity for trained compared with nontrained objects (see Supplementary Fig. 6). In summary, training did not measurably affect the spatial distribution of selectivity, and it had no significant effect on the overall response to the trained object class.

The Novel Object Selectivity Topography Was Tolerant to Changes in Behavioral Task

We found that the spatial distribution of novel-object–evoked activity was not only stable across time and training but also stable across different task contexts. During the collection of the fMRI data described so far, the shape of the objects was irrelevant to the monkeys' task—the monkeys were rewarded only for fixating and for reporting color changes (color task; see Materials and Methods). After the 3-month interval, we also scanned each monkey during performance of the very demanding within-class shape discrimination task in which they were trained (shape discrimination task; see Materials and Methods). We found little effect of this task change on the selectivity topography. For example, an examination of the voxels with significant pairwise selectivity for at least 1 object class (the same novel-object–selective ROI as used above)

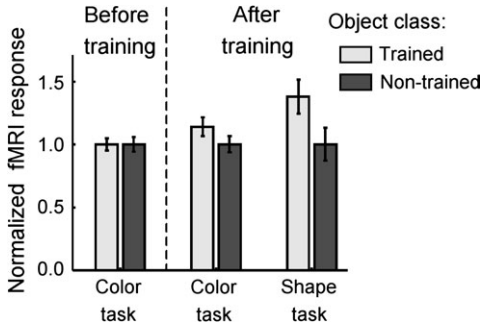

Figure 7. The effect of training on overall responsiveness in IT cortex. The results are shown for 3 experiments: color task before training (Experiment 1), color task after training (Experiment 2), and shape discrimination task after training (Experiment 3). In each monkey, the trained object category is the category for which the monkey was trained in the training phase (spikies for monkey J and smoothies for monkey M) and the untrained object category is the category for which the other monkey was trained (smoothies for monkey J and spikies for monkey M). The responses were normalized per experiment according to the overall response to the untrained object category. The overall response (expressed as PSC relative to fixation) in the different experiments was 0.90 (Experiment 1), 0.87 (Experiment 2), and 0.77 (Experiment 3). Error bars show the standard error of the mean across all time series obtained in the 2 monkeys ($N = 53$, 61, and 61 in Experiments 1, 2, and 3, respectively).

revealed nearly the same magnitude of selectivity during the shape discrimination task (31%) as already described in the data obtained during the color task (see Fig. 4*b*). This selectivity was highly significant in each monkey (J: $P < 0.01$; M: $P < 0.005$) and in none of the monkeys was it significantly different from the selectivity observed during performance of the color task before the 3-month interval (J: $P > 0.5$; M: $P > 0.5$). The correlation in pairwise selectivity between the color task and the shape discrimination task across all voxels of the standard-object–selective ROI was 0.33 averaged across monkeys and object pairs (Fig. 6). This correlation was significantly positive in each monkey (J: $P < 0.001$; M: $P < 0.05$). Although it tended to be smaller than the maximum correlation of 0.42 that is expected given the variability in the data, this effect was not significant in either monkey (J: $P > 0.4$; M: $P > 0.1$). Thus, the topography of selectivity noted in a color task was largely replicated in a different task context.

The small difference in responsivity to trained versus nontrained objects that was not significant in IT cortex during the color task (described above) was more apparent during performance of the shape discrimination task (Fig. 7)—monkey J: $P = 0.039$; M: $P = 0.099$; taking all data of the 2 monkeys together: $P = 0.010$. The size of this training-related response increase was not significantly different between the color task and the shape discrimination task ($P = 0.20$ with all data of the 2 monkeys together). This response increase is visible in Figure 6*a*, where the color associated with the trained object class (green and red for monkey J and M, respectively) tends to dominate the color maps obtained during the shape discrimination task, much more than in the color maps obtained prior to training during the color task (compare top row and third row of Fig. 6*a*). In sum, the combination of training to discriminate among objects within 1 object class and performance of that shape discrimination task resulted in somewhat higher fMRI responses to the trained class relative to the other, nontrained classes, while the spatial distribution of selectivity was largely unaffected by these factors.

The Novel Object Selectivity Topography Was Tolerant to Changes in Object Position

Although the previous analyses show that different object classes produce reliable, stable patterns of IT selectivity in at least 2 tasks, they leave open many important questions about what aspects of the object classes produce those patterns. A possible candidate is the retinotopic envelope of the 3 object classes. Even though we jittered the stimulus position, each object category was still associated with a specific retinotopic envelope (see Materials and Methods and Supplementary Figs 1 and 2). The potential importance of retinotopic envelope was confirmed by positive spatial correlations in pairwise selectivity in early retinotopic visual cortex (parafoveal V1 ROI) across the 3-month time interval (average correlation of 0.41, significant in each monkey, J: $P < 0.001$, M: $P < 0.001$).

Because neurophysiologic studies have shown that IT neurons show some tolerance to changes in object position (Kobatake and Tanaka 1994; Ito et al. 1995; Op de Beeck and Vogels 2000), we investigated the effect of changes in object position on the patterns of novel object-class selectivity uncovered here. Specifically, we attempted to isolate any position-tolerant component of the object-class selectivity by measuring the selectivity topography not only when objects were at the fovea with the original small scatter in object position (as already described) but also in separate scanning blocks in which we more strongly varied object retinal position.

In a first test, we presented the stimuli at eccentric visual field positions so that they were completely nonoverlapping with the visual field positions covered during the foveal scan sessions (specifically, objects were centered at 8.5-degree eccentricity in the 4 visual quadrants). In each monkey, we found that this change in stimulus position greatly reduced responses across ventral visual cortex compared with a control scan with foveal stimuli. In monkey J, the primary visual cortex response (parafoveal V1 ROI, see Materials and Methods) to these nonfoveal objects was 11% of the response to foveal objects, and the IT response was 38% of the response to foveal objects. In monkey M, the same position test had an even stronger effect on responsiveness (reduction to 9% and 26% of the response to foveal stimuli in V1 and in IT cortex, respectively). This strong effect of eccentricity on response strength supports findings in the literature that processing in ventral visual cortex is strongly biased toward foveal stimuli (Op de Beeck and Vogels 2000; Brewer et al. 2002). Given the substantially weaker responses to nonfoveal objects, it was not surprising that the pattern of selectivity in IT cortex (across all voxels in the standard-object–selective ROI) for 8.5-degree eccentric stimuli was only weakly reliable in monkey J (reliability of 0.29; see Materials and Methods), and not reliable in monkey M (reliability of 0.07). Because of this result, we tested the effect of changes in object position in monkey M by using jittered foveal stimuli in which the retinotopic envelope for each object class created by that jitter was altered and broadened compared with the small jitter used to obtain our original data (see Supplementary Fig. 1). Simulations confirmed that this change in stimulus position would completely alter the pattern of response in a retinotopic map (see Supplementary Material and Supplementary Figs 1 and 2). Thus, the following analyses were applied to data obtained with 8.5-degree eccentric stimuli in monkey J and to more foveal but highly position-jittered stimuli in monkey M.

Examination of IT voxels with significant pairwise selectivity for at least 1 object class using the original object position (the same novel-object–selective ROI as used above) revealed nearly the same selectivity index in blocks in which object position was varied (object-class selectivity index of 36%, see Fig. 4*b*). This selectivity was significantly different from 0 in each monkey (J: $P < 0.01$; M: $P < 0.05$) and in neither monkey was it significantly different from the selectivity in the originally tested (foveal) position (J: $P > 0.5$; M: $P > 0.35$).

In IT cortex (standard-object–selective ROI), the across-voxel correlation in pairwise selectivity computed from objects presented in the original position and computed from objects presented at the new positions was 0.44 averaged across monkeys and object pairs (Fig. 6). This correlation was significantly positive in each monkey (J: $P < 0.001$; M: $P < 0.001$) and not significantly smaller than the maximum correlation of 0.49 expected given the variability in the data (J: $P > 0.4$; M: $P > 0.30$). For comparison, we also performed the same analysis on V1 voxels (parafoveal V1 ROI) and found that the correlation in pairwise selectivity between the original position and the new position was only 0.04 ($P > 0.3$ in each monkey). The same result was found in early visual cortex as a whole—that is, all the visually active cortex ventral to the STS and more posterior than (not including) the estimated location of TEO (monkey J: 2450 voxels; monkey M: 6470 voxels). In this

large region of "pre-IT" cortex, the initial pattern of selectivity as found with the original object position was replicated after a 3-month interval when using the same object position (monkey J: $r = 0.27$, $P < 0.001$; monkey M: $r = 0.46$, $P < 0.0001$) but not when stimulus position was changed (monkey J: $r = 0.03$, $P > 0.30$; monkey M: $r = 0.09$, $P > 0.15$). This result is illustrated in Figure 8. Unlike the largely stable pattern seen in IT, the strong effect of the object position changes on the observed object-class selectivity pattern in early visual cortex is fully consistent with the differences in retinotopic envelope between object classes and experiments and their expected effect on the activation of retinotopic maps (as simulated in Supplementary Material and Supplementary Fig. 2).

Thus, the tested changes in object position largely preserved the spatial distribution of object-class selectivity in IT cortex, but not in retinotopic cortex. This indicates that the object-class selectivity that we observed in IT cortex was largely tolerant to changes in the retinal position of the objects, at least as long as the objects were not presented too eccentrically (in which case much lower responses were noted in each monkey and even a disappearance of the selectivity topography in 1 monkey).

The Novel Object Selectivity Topography Was Altered by Changes in Object Shape

Taken together, the data presented thus far are most consistent with the hypothesis that the patterns of object-class selectivity in IT reflect spatially varying preferences for what is broadly referred to as object "shape." Although these data are not by themselves sufficient to determine exactly what aspects of shape are represented, the fact that shape selectivity is found for novel objects opens the possibility of investigating the determinants of this selectivity in a more systematic way than has been done before with familiar object classes. To both confirm that shape is an important factor (unlike time, task, training, and retinal position) and to take a first step in this direction, we tested simple variations in object shape. In particular, we created new object classes (Fig. 9*a*) that were derived by altering the original object classes in the combination of global shape features (including aspect ratio) and local shape features (e.g., whether a object contains many straight lines with right corners, many spikes ending into acute angles, or many smoothly changing curves). We scanned the same 2 monkeys using these 3 new object classes to determine the impact of these shape manipulations on the patterns of selectivity described above. We merely use the labels "global" and "local" as summary labels for a series of shape changes, and we do not claim that these are the most important dimensions of shape. Instead, our main goal was to show that the selectivity topography and our methods are sensitive to changes in shape.

We found that manipulation of global and local shape features produced systematic alterations in the patterns of IT shape selectivity. As in the analyses described above, we computed the across-voxel correlation in pairwise selectivity, here taking pairs of objects that corresponded in either global or local shape. For example, we correlated the pattern of responses in the [smoothies – spikies] contrast with the [spiky smoothies – cuby spikies] contrast (corresponding global shape) and with the [smoothy cubies – spiky smoothies] contrast (corresponding local shape). Correlations tended to be positive for both global shape ($r = 0.29$) and local shape ($r = 0.13$) comparisons (Fig. 9*b*). Although these correlations tended to be significant in the individual monkeys (global properties: J: $P < 0.001$, M: $P < 0.03$; local properties: J: $P < 0.10$; M: $P < 0.03$), they were often significantly smaller than the maximum correlation of 0.55 that can be expected given the variability in the data (global shape: J: $P > 0.4$, M: $P < 0.001$; local shape: J: $P < 0.001$, M: $P < 0.001$). Thus, in contrast to our previous manipulations of time, training, and object position, these shape manipulations had a clear impact on the spatial patterns of IT selectivity. Moreover, neither of these 2 shape properties was able to fully explain the original selectivity pattern, suggesting that the underlying IT functional organization reflects both global and local shape features.

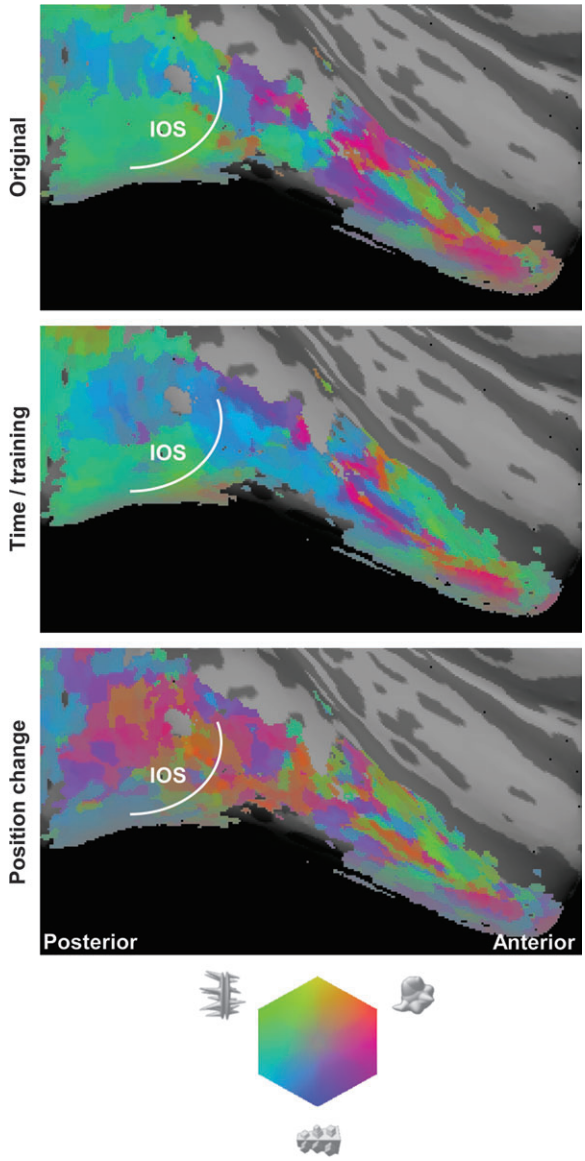

Figure 8. The effect of object position on the pattern of selectivity across the ventral visual pathway in monkey J. The pattern of selectivity in anterior regions (especially IT) as found with foveal stimuli (top 2 panels) was largely replicated with peripheral stimuli (bottom panel). In contrast, the pattern of selectivity in posterior regions was strongly altered by changes in object position. In fact, reliability analyses failed to show within-position replicability of the pattern of selectivity in posterior regions with peripheral stimuli that appeared in multiple positions (in contrast to all other selectivity maps shown in this paper). That is, the color map in posterior regions in the bottom panel was unreliable. IOS, inferior occipital sulcus.

The Relationship between the Selectivity Topography Uncovered with Novel Objects and Face Patches

Because previous fMRI studies in both humans and monkeys have revealed class-selective responses for familiar object categories (Spiridon and Kanwisher 2002; Tsao et al. 2003; Pinsk et al. 2005), we wondered how the selectivity topography uncovered with novel objects used here might be related to those earlier results. For example, if previous studies using face stimuli were uncovering selectivity that reflects the shape properties of faces, in addition to familiarity or meaning, then one might expect some relationship between the spatial patterns of object selectivity reported here and the spatial distribution of face patches, for example, a correlation across voxels between the preference for faces over other objects and the preference for one of the novel object classes.

To examine this relationship, we first replicated previously reported fMRI results of face-selective patches in an around monkey IT cortex (Tsao et al. 2003; Pinsk et al. 2005) by comparing 3 stimulus conditions: rhesus monkey faces, well-known objects (mainly fruits, toys, and lab apparel), and Fourier-scrambled images of these objects. Figure 10*a* shows color maps of the functional organization the same region of IT cortex for these 3 stimulus conditions. The face "patches" in these images are seen as red/orange regions in these maps, and their location is consistent with previous work (Tsao et al. 2003). The pattern of face selectivity across IT cortex was very reliable across time series, with a reliability (see Materials and Methods) of 0.70 and 0.85 in monkey J and M, respectively. This is higher than the reliability of the pairwise comparisons between novel object classes (0.57 and 0.74 in monkey J and M).

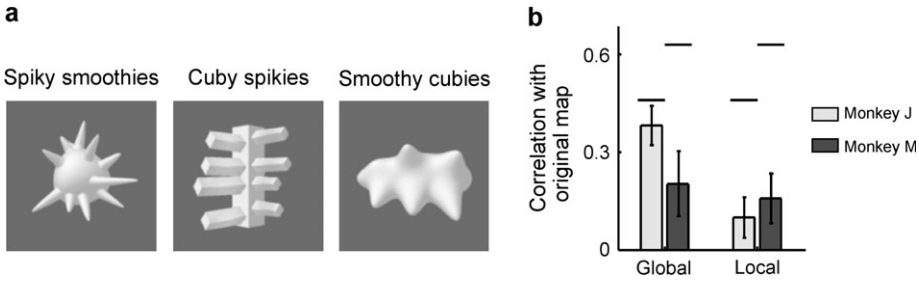

Figure 9. Changes in stimulus shape alter the selectivity topography in IT cortex. (*a*) One exemplar from each of the 3 new object classes that were used in this experiment. (*b*) The correlation of the spatial pattern of selectivity across voxels in the original measurement (Fig. 6) with the selectivity for object classes that have similar global shape properties (global) or similar local shape properties (local). The horizontal lines indicate the expected correlation if the spatial patterns were identical (as in Fig. 6*b*).

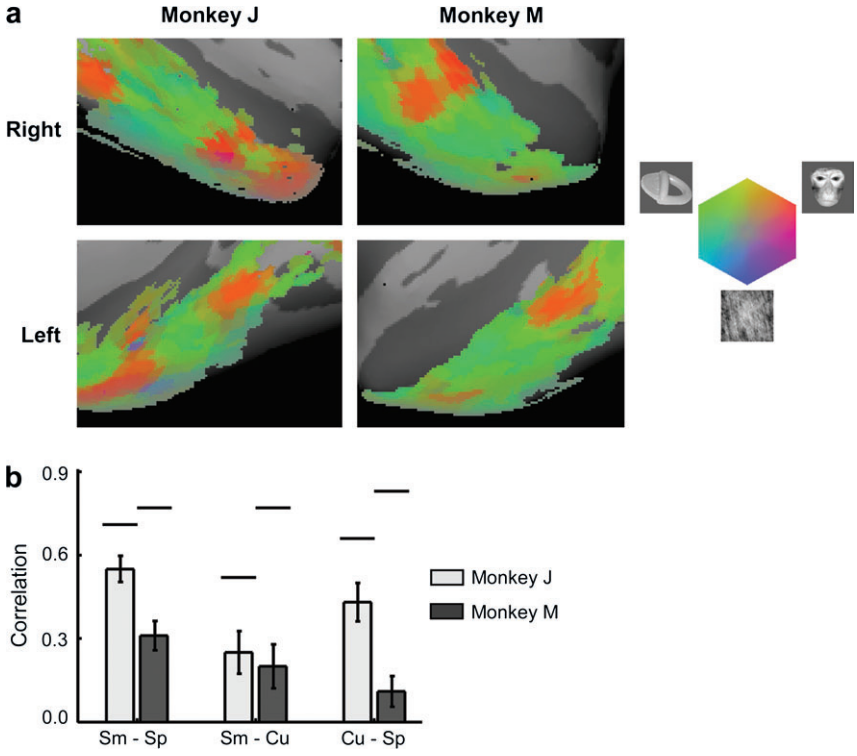

Figure 10. Selectivity maps across IT cortex for familiar objects. (*a*) Selectivity maps obtained using faces, natural objects, and scrambled images. (*b*) Correlation of the spatial distribution of selectivity for [faces − objects] with the spatial distribution of selectivity for 3 pairs of novel object classes. The horizontal lines indicate the expected correlation if the spatial pattern of face selectivity was identical to the indicated novel object selectivity pattern (as in Fig. 6*b*). Error bars show the standard error of the mean across time series. Sm, smoothies; Sp, spikies; Cu, cubies.

A qualitative comparison of the spatial location of the face patches in each monkey with the novel-object selectivity topography (Fig. 6*a*) reveals a relationship between the 2 selectivity patterns. In particular, the face patches tended to be found in or around regions that prefer smoothies (and to a lesser extent cubies) above spikies (pink regions on the original maps). To quantify this relationship, we computed a pairwise selectivity index for faces, [faces – objects], and we correlated this index with the 3 pairwise selectivity indices that were introduced before: [smoothies – spikies], [smoothies – cubies], and [cubies – spikies]. The absolute value of most of these spatial correlations was significantly above 0 ($P < 0.05$; the only exception was the correlation with [cubies – spikies] in monkey M with $P < 0.10$), with an average correlation of 0.31 across object pairs and across monkeys (see Fig. 10*b*). This analysis confirms that there is at least a partial relationship between the spatial distribution of novel objects selectivity and the distribution of face selectivity. However, each of these correlations was also significantly smaller ($P < 0.01$) than the maximum correlation (0.71 on average) that can be expected given the high reliability of the data. Thus, a substantial proportion of the distribution of novel object selectivity cannot be simply reduced to the distribution of face selectivity (and vice versa).

This conclusion was confirmed when we predicted the selectivity for novel objects after a 3-month interval from the originally measured novel object selectivity and from face selectivity. We performed a multiple regression analysis with the selectivity in each voxel in a pairwise comparison after the interval as the dependent variable and 3 predictors: the selectivity in the same object pair before the interval, face selectivity, and a constant to estimate the offset. This analysis was done for each pairwise comparison in each monkey ($n = 6$), and each time we found a larger standardized regression coefficient for the novel object selectivity (average coefficient 0.49; different from 0 across the 6 analyses: $P = 0.0002$) than for the face selectivity (average coefficient 0.12; not different from 0 across the 6 analyses: $P = 0.12$). This difference in predictive value between novel object selectivity and face selectivity was significant across voxels in each analysis ($P < 0.001$), and it was significant across the 6 analyses ($t(5) = 3.99$, $P = 0.010$). Thus, most of the temporally stable selectivity topography for novel objects could not be reduced to the distribution of face selectivity.

Finally, we found very replicable patterns of selectivity for novel objects across regions of IT cortex that showed no preference for faces over objects. The aforementioned correlation of the novel object selectivity over the 3-month time interval (average correlation of 0.52) was obtained using the data of all object-selective voxels in anterior IT cortex (standard-object–selective ROI). If the reproducible novel object selectivity reflects more than face selectivity, then a sizable correlation should be obtained if the voxels contained in the face patches are eliminated from the previous analyses. Indeed, even after removing all IT voxels with any face selectivity (even if nonsignificant, 19% of standard-object–selective ROI removed), the selectivity topography in the remaining voxels was still highly correlated across the 3-month interval (0.47 on average; significantly higher than 0 for each object pair, $P < 0.05$).

In sum, the novel object classes revealed reproducible functional organization in IT outside of the face patches, and similarly, the distribution of face patches could not fully explain the pairwise selectivity across any of the novel object classes. However, the partial relationship between the selectivity topography for novel object classes and the spatial distribution of face patches suggests that at least part of the latter is due to shape properties of faces that are also found in initially novel objects.

Discussion

To test for selectivity in IT cortex for novel object classes, we scanned 2 monkeys with fMRI while they viewed 3 classes of initially novel object stimuli. We report here the existence of a large-scale (mm+) topography for novel objects spanning IT cortex. Only a subset of voxels in IT cortex showed significant selectivity when considered individually, and this selectivity was not absolute: the response to the nonpreferred object class was about two thirds of the response to the preferred object class. However, using multivariate analyses that consider the pattern of response across multiple voxels (Haxby et al. 2001), we found a reliable topography of selectivity spanning a much larger region of IT cortex. This topography of selectivity is largely stable across time, training, changes in task, and changes in retinal position. The overall response magnitude to the different object classes is influenced by some of these manipulations, but the spatial distribution of selectivity across IT cortex is largely stable. In contrast, the topography of selectivity is sensitive to changes in object shape. Finally, this topography is not reducible to spatially organized face selectivity. Taken together, these results suggest that IT cortex contains a large-scale spatial organization for some dimensions of shape, which we will call a "shape map."

The finding of a large-scale shape map in IT cortex is largely consistent with, but not predicted by, prior work on IT cortex. Although single-unit recording and optical imaging studies have previously reported clustering of selectivity in IT cortex, the scale of that "columnar" organization (Tanaka 2003) is on the order of hundreds of microns (Fujita et al. 1992; Tanigawa et al. 2005), and such studies provide a restricted field of view of the cortex (revealing only a small part of IT cortex). Conversely, although fMRI studies in monkeys and humans have described large patches of selectivity in the cortex for specific object categories (e.g., faces, places, and bodies) (Kanwisher et al. 1997; Epstein and Kanwisher 1998; Downing et al. 2001, 2005; Tsao et al. 2003; Pinsk et al. 2005), and it has even been argued that the ventral pathway in humans contains an "object form topology" (Haxby et al. 2000), those studies have mostly used familiar object classes which confound shape with multiple other dimensions (evolutionary significance, familiarity, semantic attributes, types of processing, and associated eccentricity biases). Here we show that even novel object classes, which minimize those confounds, are associated with distinct patterns of selectivity across IT cortex.

Perhaps, the most surprising property of the shape map is its stability. Our data indicate that the shape map was present the first day that the monkeys viewed the objects, and the topography did not change for the remainder of our experiments. We made a concerted effort to induce and find changes, by scanning over several months, by extensively training the monkeys with 1 class of novel objects to change the significance of that object to the monkey, and by changing the subject's behavioral task. None of these manipulations produced a discernible effect on the object-class selectivity topography, even though our data were powerful enough to

easily detect changes resulting from partial manipulation of object shape (Fig. 9). The only effect of training that we uncovered with our methods was a small increase in the magnitude of response to the trained object class. This effect was only significant in a task context in which monkeys were actively processing shape, so the response increase might be related to attention.

Effects of Training in Monkeys versus Humans

The stability of the shape map across training stands in contrast with the outcome of our recent experiments with human subjects using the same object classes (Op de Beeck et al. 2006). Like the results described here, the human studies revealed selectivity for the novel object classes and training-associated increases in response to those objects in object-selective cortex. However, 10 days of training changed the spatial pattern of selectivity in humans. In humans, the results could be described best as a mixture of the hypotheses in *Figure 1c,d*: Part of the pretraining topography of selectivity remained in place, but part of it changed as a consequence of training. In contrast, here we found that months of training did not change the pattern in monkeys.

We can only speculate why the shape map is more stable in monkeys than humans. Perhaps, object representations in humans are more flexible than object representations in monkeys. A close comparison of the procedures and results in the 2 studies suggests other possible explanations for this difference. A first possibility is the general phenomenon that humans achieve higher performance levels on tasks that involve a comparison of multiple stimuli, like the shape task that we used here (Op de Beeck et al. 2003a). Each of the monkeys in the present study showed a very strong behavioral improvement during the training period, starting with chance performance for the largest stimulus differences in the trained object class and ending with reasonably good performance for intermediate stimulus differences (see Behavioral Tasks). Nevertheless, human subjects achieved far higher performance levels after only 10 h of training than monkeys after 3 months of training. Perhaps, the smaller effects on spatial organization of object selectivity in monkeys are linked to this behavioral difference.

A second difference between the human and monkey studies is that because the monkey studies include fewer subjects and effects must be demonstrated in each individual, the number of visual stimulus exposures prior to training was at least 5 times as large in monkeys compared with humans (more pretraining scan sessions and more stimulus exposures per session). We cannot exclude the possibility that this additional pretraining exposure somehow makes the preexisting topography of selectivity less prone to later changes in the face of training.

A third difference is that the human study used BOLD imaging and the monkey study MION imaging. However, although MION and BOLD may differ (Smirnakis et al. 2007), previous studies at 3 Tesla have shown similar patterns of activation in temporal cortex from the 2 methods (Leite et al. 2002; Tsao et al. 2003). Critically, at 3-Tesla MION is more sensitive than BOLD, arguing against the idea that we simply lacked the power to see training-induced changes in monkeys.

Is the lack of a training effect on the fMRI-determined IT selectivity topography consistent with single-unit neurophysiology studies? In general, such studies have revealed effects of experience on the overall responsiveness and selectivity in IT cortex (Kobatake et al. 1998; Baker et al. 2002; Sigala and Logothetis 2002; Freedman et al. 2006; Op de Beeck et al. 2007). Although potentially important and still underexplored, such effects at the single-unit level are typically small in magnitude (small changes in spikes/s relative to stimulus-driven changes). Training-related effects are also variable in direction, with some studies reporting training-related decreases in responsiveness (Baker et al. 2002; Freedman et al. 2006; Op de Beeck et al. 2007) and others reporting increases in responsiveness (Logothetis et al. 1995; Kobatake et al. 1998). Thus, our finding of a small training-related increase in fMRI-determined responsiveness is not obviously at odds with the single-unit literature. Although single-unit studies have not investigated potential effects of training on the spatial distribution of selectivity across IT cortex, slight increases in very local spatial clustering of similarly selective neurons have been reported in neighboring perirhinal cortex (Erickson et al. 2000). Our study is the first to look for such effects in monkey IT cortex, and we were unable to find significant effects of training on the spatial pattern of selectivity at mm+ scale. Our failure to find such effects does not rule out the possibility that training effects may occur under different training conditions, nor does it rule out the possibility that the training in our study changed the selectivity of single IT neurons without a strong change in overall fMRI-determined responsiveness or spatial pattern.

Shape Maps and Single Neurons

The stability of the shape map across the various shape-preserving manipulations agrees with the reported responses of single units in monkey IT cortex that show some tolerance for stimulus position (Ito et al. 1995; Op de Beeck and Vogels 2000) and task context (Richmond and Sato 1987; Vogels and Orban 1994). Indeed, our finding of relative stability of the object selectivity topography in IT in spite of changes in object position shows that this topography is primarily driven by differences in shape (i.e., a shape map) and not by differences in absolute retinotopic envelope. Nevertheless, our results do not rule out the possibility that monkey IT cortex also contains some kind of retinotopic map, as suggested for parts of human object-selective cortex (Hasson et al. 2002; Brewer et al. 2005; Larsson and Heeger 2006). Our data are not directly relevant for that question, although our results confirm the well-known finding that responses in IT cortex are far stronger for foveal than for peripheral stimuli.

Further work is needed to understand the relationship between the shape map revealed by fMRI and the underlying neuronal activity, especially synaptic input (e.g., as measured through local field potentials [LFPs]) and spiking activity (Logothetis 2002). In early visual cortex, hemodynamic changes measured by fMRI correlate somewhat better with LFPs than with average spiking activity (Logothetis et al. 2001). In IT cortex, average spiking activity over 3–6 mm regions is a reasonable, but far from perfect, predictor of LFP shape preferences at the center of such regions (Kreiman et al. 2006). Furthermore, LFP shape preferences correlate for sites that are up to 8 mm apart. Thus, the spatial distribution of selectivity in LFPs has a scale that is very similar to the scale of our fMRI-determined selectivity maps (see Supplementary Material and Supplementary Fig. 5). A reasonable hypothesis is that the fMRI-determined selectivity map correlates well with both an LFP-determined selectivity map and a spatially smoothed spiking activity map. Another nonexclusive hypothesis is that small pockets of very high spike selectivity across even novel object

classes underlie the fMRI-determined maps as has been reported recently for faces (Tsao et al. 2006).

The relationship between fMRI maps, cortical input (as measured with LFPs), and output (spiking activity) might also explain why we did not see any effect of training on the fMRI maps. Training might leave the input to IT cortex unaltered, and it might mostly affect how this input is processed. This hypothesis would predict a smaller effect of training on fMRI maps or LFP signals than on spiking activity. Further investigations with multiple methods of these input/output relationships and how they change as a consequence of experience might help illuminate a fundamental question: What is the functional significance of the spatial organization of shape selectivity in IT cortex? In particular, do the strong clustering of neurons in face patches (Tsao et al. 2006) or the shape topography (described here) have any computational or mechanistic benefits for the coding of faces and other objects?

Selectivity for Novel and Familiar Objects

What is the relationship between the stable, large-scale IT shape map inferred from our current data, and the previously described face patches within IT? Although face patches bear a clear relationship to the selectivity topography we found for our novel objects, this selectivity topography is not reducible to the face patches—when we remove from our analysis all voxels with any preference for faces, the correlation between repeated measurements of the shape map with our novel object classes remains intact. Although we do not have a quantitative measurement of the amount of shape similarity between our novel objects and faces, the partial similarity we observed between the pattern of selectivity among novel objects and the pattern among faces versus objects (Results, Fig. 10) might be related to a partial overlap in shape features like how compact or round a stimulus is. This is consistent with the observation of Tsao et al. (2006) that the few nonface stimuli that sometimes elicited a small response from face cells tended to have round shapes (e.g., an apple). Indeed, this partial relationship between the face patches and the shape topography revealed here raises further intriguing questions. Is a large-scale general shape map formed first in development, in turn determining where face patches appear in that map (similar to the proposed role of retinotopic biases, see Levy et al. 2001; Hasson et al. 2002)? Or might a more general shape map develop through extensive exposure to a few prominent natural object categories such as faces? A further question is whether category-selective responses found with familiar categories such as faces and bodies reflect only the perceived shape of the stimuli, or also more experience-dependent associations, for example, a hand and a leg are both body parts but have a very different shape (Erickson et al. 2000; Hanson et al. 2004; Pietrini et al. 2004).

Conclusion

Now that we have clear evidence of a stable, large-scale shape map in IT cortex, there are many open questions that will require years of investigation to answer. We do not know the number of shape dimensions coded in the shape map (Cutzu and Edelman 1996), which shape aspects are coded and which are not (Kayaert et al. 2003; Op de Beeck et al. 2003b), and the degree to which these characteristics of the shape map adhere to the predictions of current models of object recognition (e.g., O'Toole et al. 2005; Riesenhuber and Poggio 1999; Kersten et al. 2004) or to neuronal data (Kayaert et al. 2005; Kiani et al. 2007). Nor do we know whether shape maps are similar across individuals—an investigation that will require a much larger number of subjects. The present study provides a foothold that will enable us to investigate fundamental questions about the nature of object representations, the development of cortical maps, and the mechanisms underlying our remarkable ability for visual object recognition.

Funding

The McGovern Institute for Brain Research; the Pew Charitable Trusts (UCSF 2893sc to J.J.D.); the National Eye Institute (EY13455 to N.G.K.); the National Center for Research Resources (P41-RR14075, R01 RR16594-01A1, and the National Center for Research Resources Biomedical Informatics Research Network Morphometric Project BIRN002); the Mental Illness and Neuroscience Discovery Institute; the Human Frontier Science Program and the Fund for Scientific Research—Flanders (H.P.O.).

Supplementary Material

Supplementary Figures 1–3, 5, and 6 and other Supplementary Materials can be found at: http://www.cercor.oxfordjournals.org/.

Notes

We thank S. Dang, N. Knouf, J. Mandeville, R. Marini, and L. Wald for technical support, P. Aparicio and A. Papanastassiou for help with data acquisition, and C. Baker for helpful comments on the manuscript. *Conflict of Interest*: None declared.

Address correspondence to Hans P. Op de Beeck, Laboratory of Experimental Psychology, K.U.Leuven, Tiensestraat 102, B3000 Leuven, Belgium. Email: hans.opdebeeck@psy.kuleuven.be.

References

Baker CI, Behrmann M, Olson CR. 2002. Impact of learning on representation of parts and wholes in monkey inferotemporal cortex. Nat Neurosci. 5:1210–1216.

Beauchamp MS, Lee KE, Haxby JV, Martin A. 2002. Parallel visual motion processing streams for manipulable objects and human movements. Neuron. 34:149–159.

Boussaoud D, Desimone R, Ungerleider LG. 1991. Visual topography of area TEO in the macaque. J Comp Neurol. 306:554–575.

Brewer AA, Liu J, Wade AR, Wandell BA. 2005. Visual field maps and stimulus selectivity in human ventral occipital cortex. Nat Neurosci. 8:1102–1109.

Brewer AA, Press WA, Logothetis NK, Wandell BA. 2002. Visual areas in macaque cortex measured using functional magnetic resonance imaging. J Neurosci. 22:10416–10426.

Chao LL, Haxby JV, Martin A. 1999. Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. Nat Neurosci. 2:913–919.

Cox RW, Jesmanowicz A. 1999. Real-time 3D image registration for functional MRI. Magn Reson Med. 42:1014–1018.

Cronbach LJ. 1949. Essentials of psychological testing. New York: Harper.

Cutzu F, Edelman S. 1996. Faithful representation of similarities among three-dimensional shapes in human vision. Proc Natl Acad Sci USA. 93:12046–12050.

Dale AM, Fischl B, Sereno MI. 1999. Cortical surface-based analysis. I. Segmentation and surface reconstruction. Neuroimage. 9:179–194.

Denys K, Vanduffel W, Fize D, Nelissen K, Peuskens H, Van Essen D, Orban GA. 2004. The processing of visual shape in the cerebral cortex of human and nonhuman primates: a functional magnetic resonance imaging study. J Neurosci. 24:2551–2565.

DiCarlo JJ, Maunsell JH. 2000. Form representation in monkey inferotemporal cortex is virtually unaltered by free viewing. Nat Neurosci. 3:814–821.

Downing PE, Chan AW, Peelen MV, Dodds CM, Kanwisher N. 2005. Domain specificity in visual cortex. Cereb Cortex. 16:1453–1461.
Downing PE, Jiang Y, Shuman M, Kanwisher N. 2001. A cortical area selective for visual processing of the human body. Science. 293:2470–2473.
Epstein R, Kanwisher N. 1998. A cortical representation of the local visual environment. Nature. 392:598–601.
Erickson CA, Jagadeesh B, Desimone R. 2000. Clustering of perirhinal neurons with similar properties following visual experience in adult monkeys. Nat Neurosci. 3:1143–1148.
Fischl B, Sereno MI, Dale AM. 1999. Cortical surface-based analysis. II: inflation, flattening, and a surface-based coordinate system. Neuroimage. 9:195–207.
Fize D, Vanduffel W, Nelissen K, Denys K, Chef d'Hotel C, Faugeras O, Orban GA. 2003. The retinotopic organization of primate dorsal V4 and surrounding areas: a functional magnetic resonance imaging study in awake monkeys. J Neurosci. 23:7395–7406.
Freedman DJ, Riesenhuber M, Poggio T, Miller EK. 2006. Experience-dependent sharpening of visual shape selectivity in inferior temporal cortex. Cereb Cortex. 16:1631–1644.
Fujita I, Tanaka K, Ito M, Cheng K. 1992. Columns for visual features of objects in monkey inferotemporal cortex. Nature. 360:343–346.
Hanson SJ, Matsuka T, Haxby JV. 2004. Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: is there a "face" area? Neuroimage. 23:156–166.
Hasson U, Levy I, Behrmann M, Hendler T, Malach R. 2002. Eccentricity bias as an organizing principle for human high-order object areas. Neuron. 34:479–490.
Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P. 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science. 293:2425–2430.
Haxby JV, Ishai II, Chao LL, Ungerleider LG, Martin II. 2000. Object-form topology in the ventral temporal lobe. Response to I. Gauthier (2000). Trends Cogn Sci. 4:3–4.
Ito M, Tamura H, Fujita I, Tanaka K. 1995. Size and position invariance of neuronal responses in monkey inferotemporal cortex. J Neurophysiol. 73:218–226.
Kanwisher N. 2000. Domain specificity in face perception. Nat Neurosci. 3:759–763.
Kanwisher N, McDermott J, Chun MM. 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. J Neurosci. 17:4302–4311.
Kayaert G, Biederman I, Vogels R. 2003. Shape tuning in macaque inferior temporal cortex. J Neurosci. 23:3016–3027.
Kayaert G, Biederman I, Vogels R. 2005. Representation of regular and irregular shapes in macaque inferotemporal cortex. Cereb Cortex. 15:1308–1321.
Kersten D, Mamassian P, Yuille A. 2004. Object perception as Bayesian inference. Annu Rev Psychol. 55:271–304.
Kiani R, Esteky H, Mirpour K, Tanaka K. 2007. Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. J Neurophysiol. 97:4296–4309.
Kobatake E, Tanaka K. 1994. Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex. J Neurophysiol. 71:856–867.
Kobatake E, Wang G, Tanaka K. 1998. Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. J Neurophysiol. 80:324–330.
Kreiman G, Hung CP, Kraskov A, Quiroga RQ, Poggio T, DiCarlo JJ. 2006. Object selectivity of local field potentials and spikes in the macaque inferior temporal cortex. Neuron. 49:433–445.
Larsson J, Heeger DJ. 2006. Two retinotopic visual areas in human lateral occipital cortex. J Neurosci. 26:13128–13142.
Leite FP, Tsao D, Vanduffel W, Fize D, Sasaki Y, Wald LL, Dale AM, Kwong KK, Orban GA, Rosen BR, et al. 2002. Repeated fMRI using iron oxide contrast agent in awake, behaving macaques at 3 Tesla. Neuroimage. 16:283–294.
Levy I, Hasson U, Avidan G, Hendler T, Malach R. 2001. Center-periphery organization of human object areas. Nat Neurosci. 4:533–539.
Lewis-Beck MS, editor. 1994. Basic measurement. London: Sage Publications.
Logothetis NK. 2002. The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. Philos Trans R Soc Lond B Biol Sci. 357:1003–1037.
Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. 2001. Neurophysiological investigation of the basis of the fMRI signal. Nature. 412:150–157.
Logothetis NK, Pauls J, Poggio T. 1995. Shape representation in the inferior temporal cortex of monkeys. Curr Biol. 5:552–563.
Malach R, Levy I, Hasson U. 2002. The topography of high-order human object areas. Trends Cogn Sci. 6:176–184.
Malach R, Reppas JB, Benson RR, Kwong KK, Jiang H, Kennedy WA, Ledden PJ, Brady TJ, Rosen BR, Tootell RB. 1995. Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. Proc Natl Acad Sci USA. 92:8135–8139.
Op de Beeck H, Vogels R. 2000. Spatial sensitivity of macaque inferior temporal neurons. J Comp Neurol. 426:505–518.
Op de Beeck H, Wagemans J, Vogels R. 2003a. Asymmetries in stimulus comparisons by monkey and man. Curr Biol. 13:1803–1808.
Op de Beeck H, Wagemans J, Vogels R. 2003b. The effect of category learning on the representation of shape: dimensions can be biased but not differentiated. J Exp Psychol Gen. 132:491–511.
Op de Beeck HP, Baker CI, Dicarlo JJ, Kanwisher NG. 2006. Discrimination training alters object representations in human extrastriate cortex. J Neurosci. 26:13025–13036.
Op de Beeck HP, Wagemans J, Vogels R. 2007. Effects of perceptual learning in visual backward masking on the responses of macaque inferior temporal neurons. Neuroscience. 145:775–789.
O'Toole AJ, Jiang F, Abdi H, Haxby JV. 2005. Partially distributed representations of objects and faces in ventral temporal cortex. J Cogn Neurosci. 17:580–590.
Pietrini P, Furey ML, Ricciardi E, Gobbini MI, Wu WH, Cohen L, Guazzelli M, Haxby JV. 2004. Beyond sensory images: object-based representation in the human ventral pathway. Proc Natl Acad Sci USA. 101:5658–5663.
Pinsk MA, Desimone K, Moore T, Gross CG, Kastner S. 2005. Representations of faces and body parts in macaque temporal cortex: a functional MRI study. Proc Natl Acad Sci USA. 102:6996–7001.
Richmond BJ, Sato T. 1987. Enhancement of inferior temporal neurons during visual discrimination. J Neurophysiol. 58:1292–1306.
Riesenhuber M, Poggio T. 1999. Hierarchical models of object recognition in cortex. Nat Neurosci. 2:1019–1025.
Sigala N, Logothetis NK. 2002. Visual categorization shapes feature selectivity in the primate temporal cortex. Nature. 415:318–320.
Smirnakis SM, Schmid MC, Weber B, Tolias AS, Augath M, Logothetis NK. 2007. Spatial specificity of BOLD versus cerebral blood volume fMRI for mapping cortical organization. J Cereb Blood Flow Metab. 27:1248–1261.
Spiridon M, Kanwisher N. 2002. How distributed is visual category information in human occipito-temporal cortex? An fMRI study. Neuron. 35:1157–1165.
Tanaka K. 2003. Columns for complex visual object features in the inferotemporal cortex: clustering of cells with similar but slightly different stimulus selectivities. Cereb Cortex. 13:90–99.
Tanigawa H, Wang Q, Fujita I. 2005. Organization of horizontal axons in the inferior temporal cortex and primary visual cortex of the macaque monkey. Cereb Cortex. 15:1887–1899.
Tarr MJ, Gauthier I. 2000. FFA: a flexible fusiform area for subordinate-level visual processing automatized by expertise. Nat Neurosci. 3:764–769.
Tsao DY, Freiwald WA, Knutsen TA, Mandeville JB, Tootell RB. 2003. Faces and objects in macaque cerebral cortex. Nat Neurosci. 6:989–995.
Tsao DY, Freiwald WA, Tootell RB, Livingstone MS. 2006. A cortical region consisting entirely of face-selective cells. Science. 311:670–674.
Tsunoda K, Yamane Y, Nishizaki M, Tanifuji M. 2001. Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns. Nat Neurosci. 4:832–838.
Vanduffel W, Fize D, Mandeville JB, Nelissen K, Van Hecke P, Rosen BR, Tootell RB, Orban GA. 2001. Visual motion processing investigated

using contrast agent-enhanced fMRI in awake behaving monkeys. Neuron. 32:565–577.
Vanduffel W, Fize D, Peuskens H, Denys K, Sunaert S, Todd JT, Orban GA. 2002. Extracting 3D from motion: differences in human and monkey intraparietal cortex. Science. 298: 413–415.
Vogels R, Orban GA. 1994. Activity of inferior temporal neurons during orientation discrimination with successively presented gratings. J Neurophysiol. 71:1428–1451.
Wang G, Tanifuji M, Tanaka K. 1998. Functional architecture in monkey inferotemporal cortex revealed by *in vivo* optical imaging. Neurosci Res. 32:33–46.