

Corrigendum: High-resolution imaging reveals highly selective nonface clusters in the fusiform face area

Kalanit Grill-Spector, Rory Sayres & David Ress

Nature Neuroscience 9, 1177–1185 (2006); Published online 6 August 2006

Baker *et al.*¹ have drawn to our attention flaws in our analysis of voxel selectivity arising because noise can generate similar results and because this analysis excluded voxels whose amplitudes were negatively correlated across runs, which the Methods failed to state. In addition, Simmons *et al.*² note that our selectivity index incorrectly assigns high values to voxels with low or negative activations. A re-analysis of our data to address these concerns shows that our assertion that many voxels in the fusiform face area (FFA) show high selectivity to nonface objects was incorrect (Fig. 1). In the re-analysis, we used the standard general linear model³ to identify active voxels and cross-validated the results (Supplementary Methods online). Owing to these errors, Figures 4 and 8 in the original paper are not valid. The other conclusion remains correct: the fusiform face area is heterogeneous, in that regions of high selectivity for faces are intermingled with regions of lower selectivity (Fig. 1b). These flaws do not affect reproducibility analyses that included all FFA voxels (Fig. 3 and Supplementary Fig. 2), principal component analysis (PCA; Fig. 6), pattern analyses (Supplementary Fig. 5) or comparison of high-resolution and standard-resolution fMRI data (Fig. 7), which remain valid. We sincerely regret these errors.

Note: Supplementary information is available on the Nature Neuroscience website.

1. Baker, C.I., Hutchinson, T.L. & Kanwisher, N. *Nat. Neurosci.* **10**, 3–4 (2007).
2. Simmons, W.K., Bellgowan, P.S.F. & Martin, A. *Nat. Neurosci.* **10**, 4–5 (2007).
3. Worsley, K.J., Poline, J.B., Friston, K.J. & Evans, A.C. *Neuroimage* **6**, 305–319 (1997).

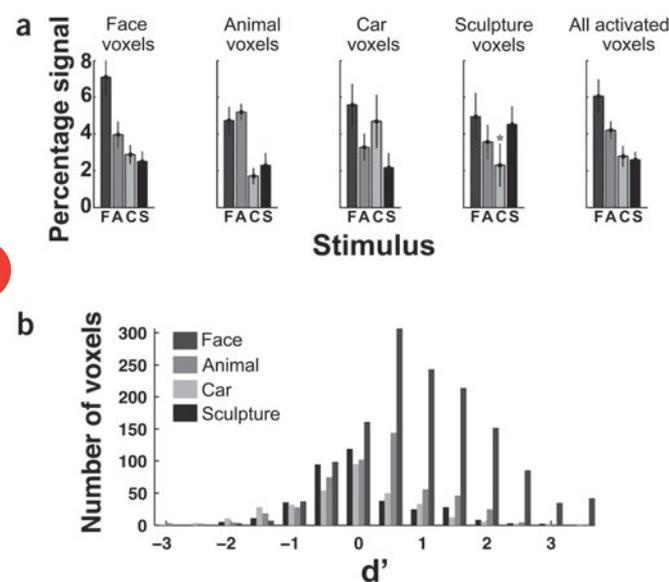


Figure 1 Re-analysis of FFA selectivity. Re-analyses were performed on 1-mm³ voxels in the FFA that were activated significantly more to objects than to the scrambled baseline (defined by a voxel-by-voxel general linear model (GLM), intact > scrambled, $P < 0.001$, no spatial smoothing). Re-analysis of average category responses with cross validation. (a) Voxels were sorted by the category that elicited the maximal response in odd numbered scans, and the average response amplitudes are plotted for even numbered scans. The rightmost column reflects the average response across all activated 1-mm³ FFA voxels. Percentage signal is measured relative to scrambled baseline. Error bars indicate s.e.m. across five subjects from published sessions. All amplitudes were significantly different from the scrambled baseline (t -test across subjects, $P < 0.05$), except for the bar marked by the asterisk. Stimuli: F, faces; A, animals; C, cars; S, sculptures. (b) Re-analysis of individual voxel selectivity using the signal detection theory measure d' , with cross validation.

$$d' = \frac{\mu_{\text{preferred}} - \mu_{\text{nonpreferred}}}{\sqrt{\frac{\sigma_{\text{preferred}}^2 + \sigma_{\text{nonpreferred}}^2}{2}}}$$

where μ and σ indicate the mean and standard deviation of responses. The preferred category for each voxel was determined as the category that yielded maximal response during odd numbered scans. d' was calculated relative to this category using data from even numbered scans.

High-resolution imaging reveals highly selective nonface clusters in the fusiform face area

Kalanit Grill-Spector^{1,2}, Rory Sayres² & David Ress³

A region in ventral human cortex (fusiform face area, FFA) thought to be important for face perception responds strongly to faces and less strongly to nonface objects. This pattern of response may reflect a uniform face-selective neural population or activity averaged across populations with heterogeneous selectivity. Using high-resolution functional magnetic resonance imaging (MRI), we found that the FFA has a reliable heterogeneous structure: localized subregions within the FFA highly selective to faces are spatially interdigitated with localized subregions highly selective to different object categories. We found a preponderance of face-selective responses in the FFA, but no difference in selectivity to faces compared to nonfaces. Thus, standard fMRI of the FFA reflects averaging of heterogeneous highly selective neural populations of differing sizes, rather than higher selectivity to faces. These results suggest that visual processing in this region is not exclusive to faces. Overall, our approach provides a framework for understanding the fine-scale structure of neural representations in the human brain.

The FFA is a region in the human brain that responds more strongly to faces than to objects¹ and is thought to be important for face perception^{1–5}. The functional role of this region is at the center of a debate about how the brain processes different types of visual stimuli and how object recognition relates to cortical organization. The domain-specific view suggests that the FFA is a cortical area specialized for processing faces^{1,6}. This view is supported by results that show that the FFA's response to faces is higher than its response to nonfaces^{1,7}, that its activation is correlated with the 'face Gestalt' rather than with low-level properties of faces^{3,4,8}, that its response correlates well with the detection and identification of faces and correlates poorly for many nonface objects⁵, and that the distributed activation across the FFA can serve to discriminate between faces and objects but not between different object categories⁹. In contrast, the distributed view suggests that the FFA is part of a distributed object-recognition system spanning the entire ventral stream^{10,11}. This view is based on other experimental results that show that the FFA's response to some nonfaces is significant¹², that the FFA shows adaptation to some nonface objects¹³, and that the distributed response across the entire ventral stream (even when excluding face-selective regions) can serve to discriminate between faces and objects¹¹.

Resolving these disparate accounts requires an understanding of the fine-scale functional organization of the FFA. This entails understanding the structure and selectivity of neural populations within the FFA and, in particular, the source of its intermediate responses to objects. We performed high-resolution fMRI (HR-fMRI, 1 mm × 1 mm × 1 mm voxels) of the FFA and asked the following questions. (i) Is the FFA a homogeneous region in which all neurons respond selectively to faces? (ii) What is the spatial distribution of regions that are selective to

faces and nonfaces across fusiform cortex? (iii) Is the selectivity to faces different from the selectivity to nonfaces?

Standard-resolution fMRI (SR-fMRI) of the FFA shows maximal responses to faces and intermediate responses to nonfaces. This pattern of response is consistent with four distinct models of the fine-scale functional organization of the FFA (**Fig. 1**). It may reflect (i) responses of a uniform neural population with broadly tuned neurons that respond maximally to faces but also respond significantly to nonfaces¹⁴ (**Fig. 1a**), (ii) averaging between a region with a high proportion of highly face-selective neurons (neurons that respond strongly to faces, but very little to nonfaces¹⁵) and a distinct region containing object-selective neurons⁷ (for example, the lateral occipital cortex (LOC)¹⁶; **Fig. 1b**), (iii) a heterogeneous structure within the FFA consisting of highly face-selective neural populations and broadly tuned neural populations responsive to both nonfaces and faces (**Fig. 1c**), or (iv) a highly heterogeneous structure within the FFA in which neural populations that are highly selective to faces are spatially interdigitated with neural populations that are highly selective to nonfaces (**Fig. 1d**). Under this model, the higher response to faces measured with SR-fMRI is due to more face-selective neurons rather than higher selectivity to faces. Therefore, the first two models (depicted in **Fig. 1a,b**) predict a homogeneous structure and specialization for face processing across the FFA, but differ in explaining the source of fMRI responses to nonfaces. In contrast, the last two models (**Fig. 1c,d**) propose a heterogeneous structure within the FFA (but differ in the selectivity for nonfaces). These models can be resolved with HR-fMRI because they yield a different pattern of high-resolution responses in terms of both spatial scale and selectivity of responses (**Fig. 1**, HR-fMRI).

¹Department of Psychology and ²Neuroscience Program, Stanford University, Stanford, California, 94305, USA. ³Department of Neuroscience, Brown University, Providence, Rhode Island, 02912, USA. Correspondence should be addressed to K.G.-S. (kalanit@psych.stanford.edu).

Received 6 June; accepted 10 July; published online 6 August 2006; doi:10.1038/nn1745

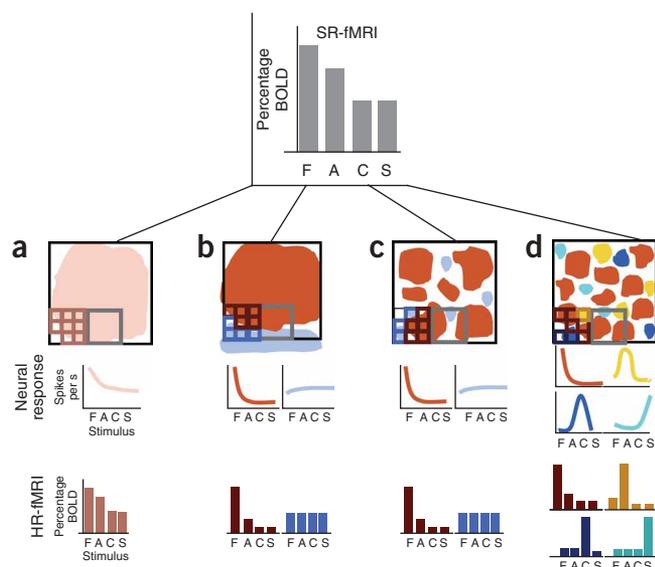


Figure 1 Models for the fine-scale functional organization of the FFA. Rows indicate spatial distribution of neurons across fusiform cortex (top), and the responses observed with single-unit electrophysiology (middle) and high-resolution fMRI (HR-fMRI; bottom). *x*-axis, stimulus (F, face; A, animal; C, car; S, sculpture). Faded colors, regions with low-selectivity neurons. Saturated colors, region with high-selectivity neurons. Red, regions containing face-selective neurons. Spatial scale is illustrated by a representative SR-fMRI voxel (gray square) and several HR-fMRI voxels (small squares). The large black square indicates the putative border of the FFA measured with SR-fMRI (side ~ 1 cm). Bar graphs indicate average fMRI responses across voxels that prefer a category. **(a)** Homogeneous broadly selective model: FFA contains broadly tuned neurons that prefer faces and also respond robustly to nonfaces. **(b)** Homogeneous highly selective model: FFA contains highly selective face neurons that respond strongly to faces but do not respond significantly to nonfaces. Blue regions indicate adjacent object-selective cortex that contains neurons that respond to objects. **(c)** Patchy model: FFA contains a mixture of face-selective and object-selective neurons. Face-selective neurons are highly selective to faces, but object-selective neurons respond to many categories. This model predicts a finer spatial scale of heterogeneous neural populations than that in **b**. **(d)** Highly heterogeneous model: FFA consists of highly selective neural populations for both faces and nonfaces that are spatially interdigitated. There is a larger number of neurons that are selective to faces than to nonfaces, but selectivity is not different between categories.

To distinguish between these models of FFA structure, we conducted HR-fMRI (voxels of $1\text{ mm} \times 1\text{ mm} \times 1\text{ mm}$) of the FFA. Five subjects were scanned using a 3-Tesla scanner equipped with a surface coil while they viewed blocks of faces and three nonface object categories (animals, cars and abstract sculptures; **Supplementary Fig. 1** online) that alternated with scrambled versions of these stimuli. To examine the generality of our results, subjects participated in eight runs using different images and block orders. Subjects repeated the experiment using SR-fMRI (voxels of $3\text{ mm} \times 3\text{ mm} \times 3\text{ mm}$) so that we could relate HR-fMRI to SR-fMRI data.

RESULTS

We defined the FFA using SR-fMRI (**Fig. 2**), and then examined the properties of this region using HR-fMRI. We measured the degree of selectivity of FFA voxels to different categories, the spatial distribution of voxels according to their preferred stimulus, and the information available in the distributed activation across the FFA (Methods). Selectivity of responses was measured using a selectivity index that reflects both the preferred category for a voxel and the degree of preference for that category. A value of 1 indicates maximal preference; 0 indicates no preference (Methods).

Standard-resolution fMRI of the FFA

The FFA was defined as the cluster of $3\text{ mm} \times 3\text{ mm} \times 3\text{ mm}$ voxels in the fusiform gyrus and occipito-temporal sulcus (OTS) that activated more strongly to faces than to objects ($P < 10^{-3}$, voxel level; **Fig. 2a,b**). The average response amplitude across the entire FFA was maximal to faces and intermediate (but significant) to objects (**Fig. 2c**; $P < 0.05$, *t*-test across subjects). We examined the selectivity of SR-FFA voxels to different categories (Methods and **Fig. 2d–e**). As shown in the

histogram (**Fig. 2d**), almost all SR-FFA voxels were face selective, with a mean selectivity of 0.4 ± 0.2 . In each subject, there were an average of 19 ± 7 (mean \pm s.d. across subjects) standard-resolution voxels that responded maximally to faces, but only a few voxels responded maximally to a nonface category (animals: 1.9 ± 2 ; cars: 0.9 ± 1.8 ; sculptures: 0.2 ± 0.4). Both face-selective and nonface-selective populations responded robustly to faces and significantly to the other nonface categories (**Fig. 2f**; $P < 0.05$, *t*-test across subjects). Even voxels that

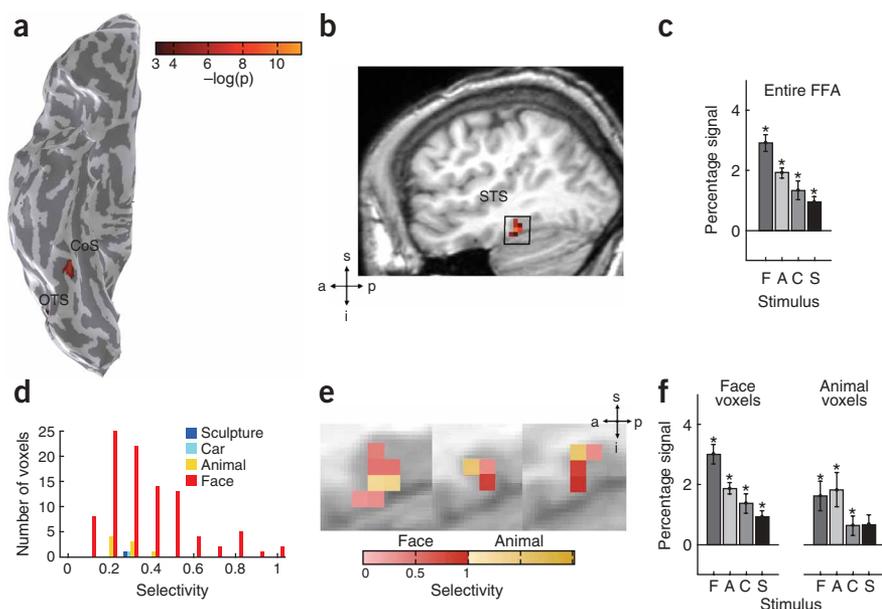


Figure 2 SR-fMRI ($3\text{ mm} \times 3\text{ mm} \times 3\text{ mm}$) of the FFA. **(a,b)** Activation maps: faces $>$ objects ($P < 10^{-3}$, voxel level) of subject 3. **(a)** Map on the inflated cortical surface; **(b)** in-plane view (pseudo-sagittal slice). Axes: a, anterior; p, posterior; s, superior; i, inferior. **(c)** Mean amplitude of response in FFA to different categories relative to the scrambled baseline averaged across five subjects. Error bars indicate s.e.m. across five subjects. $*P < 0.05$ (higher response as compared to scrambled; *t*-test across subjects). **(d)** Histogram of the selectivity index of SR-FFA voxels across five subjects. **(e)** Selectivity map of standard FFA voxels on a magnified region in the vicinity of the FFA (corresponding to the black square in **b**). Data are presented for three consecutive in-plane slices of subject 3. **(f)** Responses of face-selective and animal-selective voxels averaged across five subjects. Conventions same as in **c**.

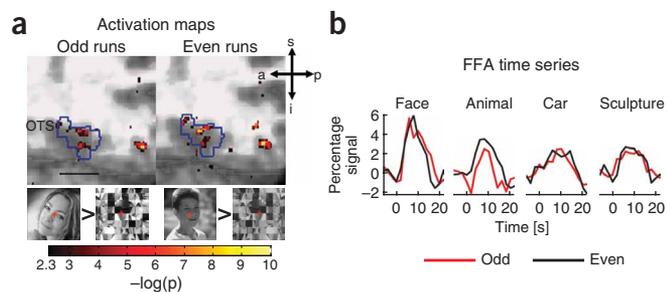


Figure 3 Reliability of HR-fMRI. Data from subject 1. (a) Voxel-by-voxel activation maps (faces > scrambled, $P < 0.005$, voxel level) for two independent data sets (odd and even runs) on in-plane slices. Blue contour indicates the boundary of the right FFA. OTS, occipito-temporal sulcus. (b) Time series averaged across the FFA separately for odd and even runs. Percentage signal is calculated versus the scrambled baseline. Time 0 indicates onset of each block.

responded maximally to nonfaces (for example, 'animal-selective') still had robust face responses that were greater than the average response to all nonfaces (that is, 'animal-selective' voxels were also 'face-selective', Fig. 2f). Therefore, all SR-FFA voxels showed a robust and preferential response to faces. Overall, at standard resolution, the FFA seems to be a uniform region consisting almost entirely of a homogeneous group of voxels that respond maximally to faces and also significantly to nonface stimuli ($P < 0.05$, t -test across subjects).

High-resolution fMRI of the FFA

We conducted HR-fMRI of the FFA in the same subjects (Methods). We found strong face responses (faces > scrambled, $P < 10^{-3}$, voxel level) in localized clusters within the FFA, but these responses did not uniformly cover the FFA (Fig. 3a). Because HR-fMRI is a new method, we first examined the reliability of HR-fMRI in each subject using three measures: spatial reliability (that is, whether activation maps from two independent data sets (odd and even runs) using different images were consistent), temporal reliability (whether time series from two independent data sets were consistent), and the functional signal-to-noise ratio of HR-fMRI.

Voxel-by-voxel activation maps for HR-fMRI were robust and reliable in each subject (Fig. 3a, one subject; Supplementary Fig. 2 online, all subjects). For all subjects and categories, there was a significant correlation (ρ) between activation maps of independent data sets (odd and even runs) versus scrambled images (ρ_{faces} : 0.43 ± 0.06 ; ρ_{animals} : 0.36 ± 0.08 ; ρ_{cars} : 0.3 ± 0.03 ; $\rho_{\text{sculptures}}$: 0.34 ± 0.06 ; mean \pm s.e.m. across subjects). All correlations were significant for individual subjects ($P < 0.01$) and across subjects (t -test, $P < 0.02$). Because different images were used in each run, this demonstrates that activation maps were reliable across runs and generalized across images of a category. Similarly, in all subjects the FFA time series during odd runs was significantly correlated with the time series during even runs (Fig. 3b, one subject; Supplementary

Fig. 2, all subjects; $\rho_{\text{time series}} = 0.66 \pm 0.05$; individual subject: $P < 10^{-6}$; across subjects: $P < 10^{-5}$). The functional signal-to-noise ratio measured across the FFA was 14.3 ± 2 dB, indicating that the signal amplitude was five times higher than the fluctuations of the scrambled baseline. These analyses indicate that we measured selective and reliable responses in both space and time at a voxel size of 1 mm^3 .

Selectivity of high-resolution FFA responses

Selectivity maps (Methods) were calculated separately for odd and even runs and were reliable ($\rho_{\text{selectivity}} = 0.89 \pm 0.05$, $P < 10^{-6}$) for each subject. Notably, analysis of the selectivity of HR-fMRI voxels revealed that the FFA contains clusters of voxels with heterogeneous selectivity to different categories: some were selective to faces and some were selective to nonfaces (Fig. 4). Clusters of high-resolution voxels selective to nonfaces were spatially intermingled with clusters of voxels that were selective to faces (Fig. 4a, one subject; Supplementary Fig. 3 online, all subjects). Both nonface- and face-selective voxel populations showed strong responses and high selectivity to the preferred category (Fig. 4b). For example, 'face voxels' responded robustly to faces but not to cars or sculptures, and 'car voxels' responded robustly to cars but not to faces, animals or sculptures.

Analysis of the distribution of selectivity indices of high-resolution FFA voxels (Fig. 4c) revealed a substantial number of voxels within the FFA selective to each of the four categories (faces, animals, cars and sculptures). Also, for each category there was a considerable set of voxels with a selectivity index of 1, indicating that the FFA contains both neural populations with high selectivity to faces and neural populations with high selectivity to specific nonface categories. Overall, high-resolution voxels showed significantly higher selectivity than standard-resolution voxels ($P < 10^{-6}$, t -test). Notably, the mean selectivity of face-preferring voxels was not higher than that for other

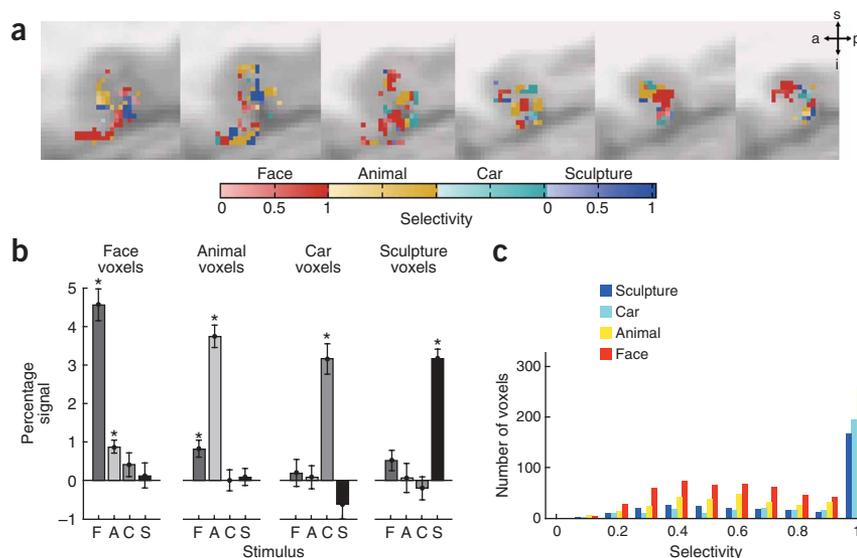


Figure 4 HR-fMRI ($1 \text{ mm} \times 1 \text{ mm} \times 1 \text{ mm}$) of the FFA. (a) FFA selectivity map (subject 3). Data are shown for a magnified region in the vicinity of the FFA, on six consecutive 1-mm in-plane slices. Color bar, preferred category and selectivity level. The leftmost slice is the most lateral. (b) Response amplitudes of voxel populations within the FFA to different stimuli. Data are averaged across five subjects. Baseline is scrambled condition. Error bars indicate s.e.m. across subjects. $*P < 0.05$ (higher response as compared to scrambled; t -test across subjects). (c) Histogram of the selectivity index of high-resolution FFA voxels across five subjects.

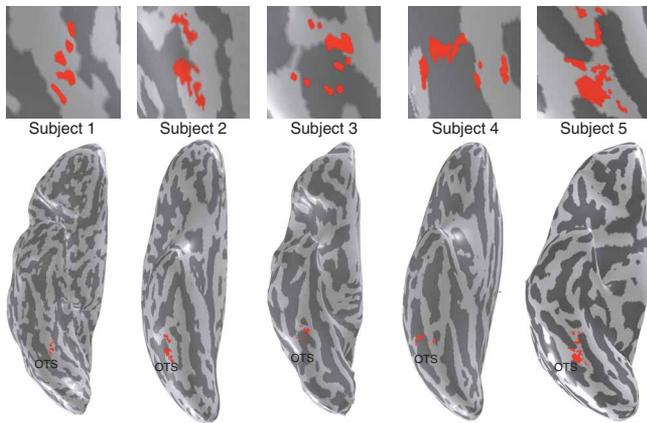


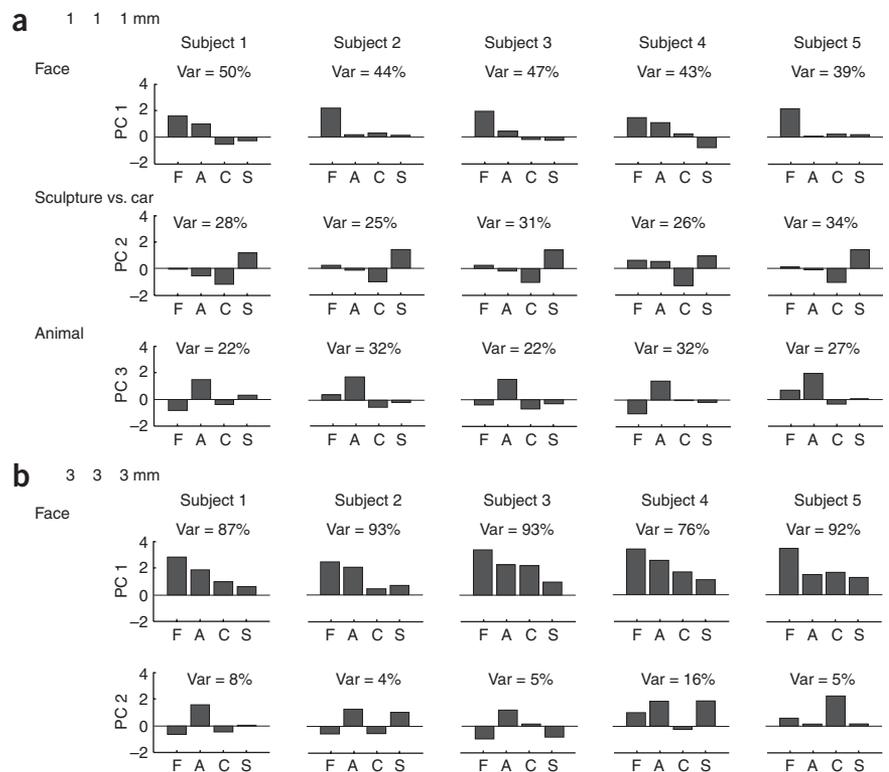
Figure 5 HR-fMRI face-selective patches in the FFA on the inflated cortical surface. Face-selective patches (red) in the FFA were determined from the selectivity map. These regions were projected on the gray matter surface of each subject. Top, magnified view of the FFA. Bottom, entire right hemisphere. OTS, occipito-temporal sulcus.

categories (selectivity_{faces}: 0.77 ± 0.25 ; selectivity_{animals}: 0.80 ± 0.24 ; selectivity_{cars}: 0.86 ± 0.23 ; selectivity_{sculptures}: 0.80 ± 0.25 ; mean \pm s.d. across subjects).

Although selectivity did not differ across categories, we found that there were significantly more high-resolution voxels that were selective to faces than to nonfaces ($P < 0.01$, t -test). The average number of high-resolution voxels in the FFA that were selective to faces was 155 ± 36 (mean \pm s.d. across subjects); for nonfaces, the number of voxels was lower (animals: 104 ± 36 ; cars: 63 ± 20 ; sculptures: 63 ± 25). We also found that the response amplitude to the preferred category was highest for face-preferring voxels (paired t -test with Bonferroni correction, $P < 0.001$).

The spatially heterogeneous structure of the FFA (Fig. 4a and Supplementary Fig. 3) was not an artifact of data analysis on in-plane slices. Projecting selective subregions within the FFA onto the cortical surface of each

Figure 6 PCA of the FFA. (a,b) The first few principal components that accounted for at least 95% of the variance in the data for both HR-fMRI and SR-fMRI. Each column shows the principal components for one subject. Var, percent variance (out of total variance) that each component accounts for. For the most part, rows are ordered by decreasing component variance. (a) HR-fMRI. The first three components reflected information about faces, sculptures versus cars, and animals, respectively. For consistency with the type of component, we plotted, for subjects 2 and 4, the third component before the second (see variance). For subjects 1 and 4, the first component contained information about both faces and animals, and the third component distinguished between animals and faces. (b) SR-fMRI. The first component showed a graded response across categories and did not separate well between them (except that responses to faces were higher overall). In some subjects (1 and 5), the second component was able to separate one additional category, but these components were not sufficient to provide linear separation between the four categories.



subject (Methods) revealed that face-selective subregions measured at high resolution were not contiguous on the cortical sheet (Fig. 5). Furthermore, the patchy structure was evident both on the anatomical volume and along the cortical surface for all categories and subjects (Supplementary Fig. 4 online), which suggests that it may be an inherent property of the functional organization in fusiform cortex.

These selectivity analyses revealed that the FFA contains a reliable heterogeneous structure consisting of subregions that respond robustly and are highly selective to faces as well as subregions that respond robustly and are highly selective to nonfaces. These results are consistent with the highly heterogeneous model of the FFA (Fig. 1d), suggesting that the FFA contains neural populations that are highly selective to faces as well as distinct neural populations that are highly selective to nonfaces.

Principal component analysis of FFA data

To complement the selectivity analyses, we performed principal component analyses (PCA; ref. 17) of both HR-fMRI and SR-fMRI data (Methods). The goal of PCA was twofold: to use a data-driven approach with minimal priors to validate the structure of our data, and to examine whether FFA voxels contain independent information about face and nonface categories. The outcomes of PCA are the principal components (PCs) along which the data covary and the fraction of the variance of the data explained by each component. We reasoned that more PCs would be required to explain the heterogeneity of HR-fMRI data than SR-fMRI data and that HR-fMRI PCs should show sharper selectivity than SR-fMRI PCs.

PCA of HR-fMRI revealed that three PCs were necessary for explaining 95% of the variance of this data (variance accounted for by the first three PCs was $45 \pm 4\%$, $31 \pm 2\%$ and $24 \pm 2\%$, respectively, mean \pm s.d. across subjects). These PCs were similar across subjects and provided informative axes (Fig. 6a): the first PC contained 'face'

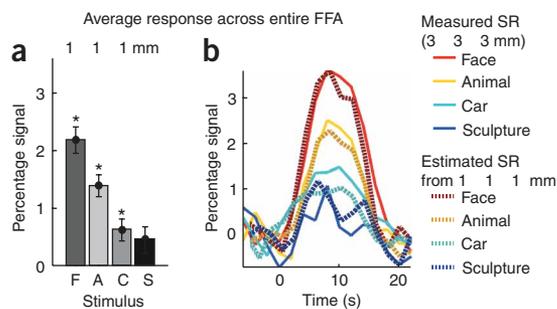


Figure 7 Comparison between HR-fMRI and SR-fMRI of the FFA. **(a)** Average response across the entire FFA measured with HR-fMRI. Response amplitudes were averaged across five subjects. Error bars indicate s.e.m. $*P < 0.05$ (compared to scrambled; *t*-test across subjects). **(b)** SR-fMRI FFA time series averaged across five subjects (solid) versus estimated SR-fMRI FFA time series (dashed) from HR-fMRI data. Signal is measured relative to scrambled. Time 0 indicates onset of block.

information, the second PC contrasted between ‘cars’ and ‘sculptures’, and the third contained ‘animal’ information. These axes demonstrated that the factors along which high-resolution voxels vary are responses to specific categories. Thus, PCA confirmed our selectivity analysis and showed that the FFA contains information about both face and nonfaces, which are linearly separable.

Indeed, an independent linear classifier trained on activation patterns from half the runs decoded with 100% success which of the four categories subjects viewed during the other half of runs (**Supplementary Note** and **Supplementary Fig. 5** online). Moreover, pattern analyses^{10,11} of split halves of our data (**Supplementary Note**) revealed that the distributed high-resolution activations across the FFA for specific categories were reliable (high on-diagonal values of the cross-correlation matrix of odd versus even runs) and activations for different categories were distinct (low off-diagonal cross-correlation values, **Supplementary Fig. 5**).

In contrast, PCA of SR-fMRI revealed that $90 \pm 4\%$ of the variance of SR-fMRI could be explained by just one principal component with graded responses to both faces and nonfaces (**Fig. 6b**). This suggests that the effective dimensionality of SR-fMRI data is 1. Therefore, from SR-fMRI, linear classifiers cannot decode the four stimulus categories. Because the responses to faces are higher, linear classifiers were able to distinguish from the distributed standard-resolution activations whether subjects viewed faces but were at chance in their ability to distinguish between nonface categories (**Supplementary Fig. 5**). PCA showed that SR-FFA voxels are rather uniform in their selectivity and the main component explaining their activation contains information about many categories.

Overall, PCA and pattern analyses confirmed the reliability of the heterogeneous FFA structure resolved by HR-fMRI and further demonstrated that there is distinct information about several categories (both face and nonface categories) in HR-FFA activations.

High-resolution fMRI explains standard fMRI data

Finally, we validated the highly heterogeneous model of the FFA (**Fig. 1d**) by examining whether heterogeneous HR-fMRI activations can account for homogeneous SR-fMRI responses. We reasoned that if SR-fMRI responses reflect averaging of heterogeneous populations, the overall response across the entire FFA would be similar across resolutions. Indeed, the average response across the entire FFA response when measured with HR-fMRI showed the same response profile as SR-fMRI (compare **Fig. 7a** and **Fig. 2c**):

it was maximal for faces and intermediate (but significant; $P < 0.05$, *t*-test across subjects) for nonfaces.

We then quantified the relation between the entire FFA response measured with SR-fMRI and heterogeneous high-resolution responses, by estimating the SR-FFA time series using a weighted sum of high-resolution category-selective time series (Methods). Weights were proportional to the fraction of high-resolution voxels that preferred each category (for example, the time series of face-selective voxels was multiplied by the fraction of high-resolution face-selective voxels in the FFA, and the time series of car-selective voxels was multiplied by the fraction of high-resolution car-selective voxels). We found a high correlation ($\rho = 0.93$, $P < 10^{-6}$) between the estimated standard-resolution time series (**Fig. 7b**, dashed) and the measured SR-fMRI FFA response (**Fig. 7b**, solid).

These results suggest that the maximal responses to faces and the intermediate responses to nonfaces in the FFA are due to averaging across heterogeneous and highly selective neural populations of different sizes. Higher FFA response to faces is due to a larger number of face-selective voxels (suggesting a larger face-selective neural population). Intermediate responses to nonfaces correspond to the lower number of voxels preferring nonfaces (rather than to a lower selectivity to nonfaces). These analyses validate the highly heterogeneous model (**Fig. 1d**) and demonstrate that the heterogeneous fine-scale structure of the FFA we found with HR-fMRI accounts for the graded responses measured with SR-fMRI.

HR-fMRI of adjacent object-selective cortex

How does high-resolution activation outside the FFA compare with that within the FFA? Is the pattern of highly selective category-specific clusters a property of a larger region in ventral temporal cortex? To examine these questions, we used SR-fMRI to define a second object-selective region using the standard criterion to localize the LOC (refs. 16, 18–20; here we used cars and sculptures $>$ scrambled, $P < 10^{-5}$, voxel level). Within the limited volume of our high-resolution prescription, we were able to localize in several subjects the anterior portion of the LOC (pFUS/OTS in previous studies^{16,19,20}), which was located posterior and lateral to the FFA (**Fig. 8a**, OTS). We then examined responses in this region using both SR- and HR-fMRI.

SR-fMRI revealed that the mean response across this region was significantly higher ($P < 0.05$; *t*-test across all subjects) for all four categories compared to that for the scrambled images, with no difference between categories (**Fig. 8a**). We observed several differences between object-selective cortex and FFA: SR-fMRI of object-selective cortex revealed that individual standard-resolution voxels contained a heterogeneous preference to different categories (**Fig. 8b**) and the selectivity to the preferred category (mean selectivity = 0.22 ± 0.05) was significantly weaker than the selectivity to faces in FFA (*t*-test, $P < 0.002$). These weak (but reliable) biases for different categories observed at standard resolution may explain why pattern analysis methods can decode the responses to different categories from the distributed response across object-selective cortex¹¹ (**Supplementary Fig. 5**).

HR-fMRI of this object-selective region revealed robust and reliable responses (**Supplementary Table 1** online, OTS). Similar to the FFA, high-resolution voxels showed high selectivity to particular categories including both faces and nonfaces (**Fig. 8c**; mean selectivity_{faces}: 0.73 ± 0.27 ; selectivity_{animals}: 0.76 ± 0.27 ; selectivity_{cars}: 0.78 ± 0.27 ; selectivity_{sculptures}: 0.8 ± 0.24). Overall, selectivities were somewhat lower than in the FFA (*t*-test, $P < 0.02$, except for sculptures $P = 0.4$). The distributed activation for different categories was reliable and distinct, and an independent classifier successfully decoded from the distributed high-resolution responses which category subjects viewed

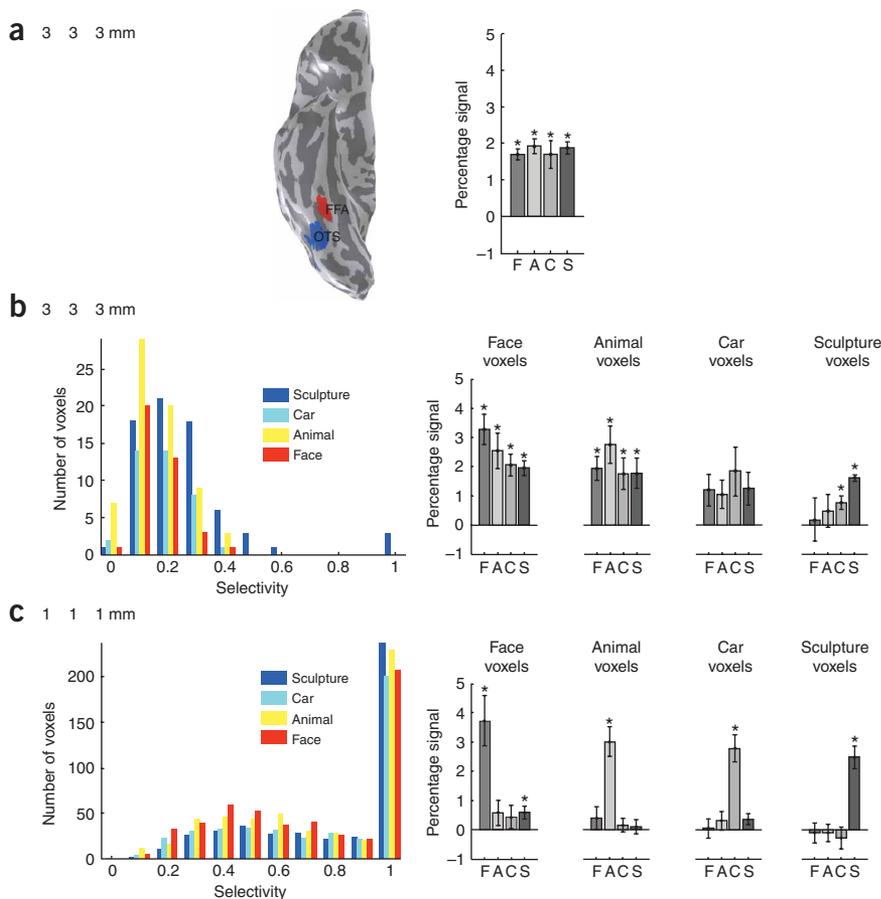


Figure 8 Responses of object-selective cortex along the occipito-temporal sulcus (OTS). **(a)** SR-fMRI. Left, location of the FFA and an object-selective ROI consisting of the anterior part of the lateral occipital complex (LOC) contained within the OTS. Right, average amplitude of OTS ROI for different stimuli across four subjects. Error bars represent s.e.m. across four subjects. $*P < 0.05$ (compared to scrambled; *t*-test across subjects). **(b)** SR-fMRI. Left, selectivity histogram of SR-fMRI voxels across four subjects. Right, mean response amplitudes compared to the scrambled baseline for different stimuli for voxel populations sorted by selectivity. Error bars indicate s.e.m. across subjects. $*P < 0.05$ (compared to scrambled; *t*-test across subjects). Conventions same as in **a**. **(c)** HR-fMRI. Same as in **b**, but for 1 mm \times 1 mm \times 1 mm data.

across heterogeneous and highly selective populations of different sizes, rather than higher selectivity to faces.

Our data provide new insights about the fine-scale structure of face-selective cortex. First, they challenge the notion that the FFA is a homogeneous area, as face-selective patches were spatially intermixed with patches that responded robustly and were highly selective to nonfaces. Second, the similar selectivity profile for faces and nonfaces at high resolution suggests that the representation of faces may not be fundamentally different from the representation of other objects. Third, the sharp selectivity profiles to specific object

categories in localized subregions in the FFA (**Fig. 4**, and **Supplementary Figs. 3,4**) and the linearly independent components of face and nonface information revealed by PCA indicates that face and object representations in the FFA do not overlap. This suggests that distinct subregions within the FFA may contribute to the perception of different object categories. Thus, previous findings that SR-fMRI FFA responses correlate with both face and animal recognition⁵ may have been driven by distinct ‘animal’ and ‘face’ subregions within the FFA.

Our results suggest two hypotheses for the functional organization in this part of cortex. First, the face- and nonface-selective subregions may be part of a common cortical region, which processes both face and nonface stimuli. Alternatively, the face-selective subregions may constitute the “true FFA” (ref. 7) (and may contain only highly selective face neurons¹⁵), whereas the other subregions may comprise a segregated subsystem. However, the fact that face-selective patches are not spatially contiguous on the cortex (**Fig. 5**) raises the question of which of them might be considered the FFA, or whether these spatially segregated subregions might behave functionally as a computational unit. Future studies may elucidate whether these face patches are interconnected, which would allow them to operate as one computational unit (for example, by studying connectivity between subregions in the FFA; ref. 22). Moreover, by combining microstimulation and fMRI (refs. 23–25), for instance, it may be possible to determine whether activation in these face patches is necessary and sufficient for face perception.

Why are face patches (and perhaps face neurons) more numerous in the FFA? One possibility is that because of their complexity and the type of information extracted from them, faces require a more

(**Supplementary Fig. 5**). Unlike the FFA, we found a similar number of high-resolution voxels that responded to different categories (number of face high-resolution voxels: 130 ± 45 ; animals: 129 ± 42 ; cars: 107 ± 32 ; sculptures: 111 ± 35 ; no significant difference between categories, *t*-test, $P > 0.7$). Because SR-fMRI responses reflect averaging over heterogeneous neural populations of similar size, the average standard-resolution response from this region does not show preference to a specific category.

Overall, these results suggest that a large object-selective region in the human ventral stream contains highly-selective neurons, global biases in standard fMRI reflect the relative size of selective neural populations, and this selectivity seems to occur on a scale of a few millimeters that can be resolved with HR-fMRI. The weaker selectivity typically observed in SR-fMRI is a consequence of its inability to resolve the small scale of selective neural populations rather than a general pattern of broad response to many stimuli.

DISCUSSION

Our results indicate that there is a reliable fine-scale structure within the FFA. This region contains a mixture of heterogeneous neural populations of voxels: some are highly selective to faces and some are highly selective to nonface objects (**Fig. 4** and **Supplementary Fig. 3**). Although the FFA contains a preponderance of face-selective voxels (suggesting a majority of face-selective neurons), it is not uniquely activated by faces, and face-selective patches do not form a continuous subregion of the FFA (**Fig. 5**). Notably, HR-fMRI showed that the maximal FFA responses to faces and the intermediate responses to nonfaces measured with SR-fMRI (refs. 12,13,21) reflect averaging

extensive representation than objects. A second possibility is that expert categories are over-represented in cortex^{26,27}. A third possibility is that this region reflects a foveal representation in which face features are preferentially represented^{28,29}. Future studies may address these possibilities by examining whether distinct face patches represent different aspects of faces such as viewpoints of faces^{30,31}, face subtypes (for example, male or female) or face-feature combinations³²; whether the emergence of face patches depends on experience^{33,34}; and whether there is a relationship between retinotopy and selectivity in these subregions²⁹.

A question that remains is whether neural responses from this region are most informative about the category of stimulus being shown or about other factors. Our research suggests that category may be one source of information. Future research is necessary for elucidating whether there is a systematic representation of category, global shape, visual features, conceptual properties and/or retinotopy in these regions. Data from our current study may give some insights. We examined whether common local features can explain these category-specific patches by comparing activations to whole and scrambled images of specific categories (**Supplementary Note**). We found that global structure rather than local features produced the robust and selective activation patterns we report here (**Supplementary Note** and **Supplementary Fig. 6** online).

Another open question is whether the FFA is distinct from the rest of object-selective cortex. Answering this question requires further research, as there is no consensus on the criteria for defining areas in the human ventral stream. Notably, within both fusiform and OTS cortex we found localized patches with high selectivity to particular categories, suggesting that this patchy organization on the few-millimeter scale may be a general organizing principle across the human ventral stream. The higher signals and selectivity of high-resolution patches (compared to that of the standard-resolution voxels) suggests that, within these patches, there may be clusters of neurons with similar selectivity^{15,30,35}. In turn, this suggests that the functional organization in higher-level visual cortex may follow principles similar to those in low-level cortex, in which neurons with similar selectivity are clustered.

What are the computational advantages of the localized and highly selective organization we report? One advantage is that this kind of encoding represents information in a way that is easier to read in subsequent levels of processing³⁶. Second, this representation may optimize the capacity of this region for storing patterns³⁶ and may provide the capacity to code the many thousands of visual categories humans can classify. Our findings that a few tens of high-resolution voxels in the FFA are sufficient to distinguish four categories (**Supplementary Fig. 5**) suggests that the entire FFA may have the capacity to code a multitude of object categories.

Implications for the interpretation of standard fMRI data

The present data explain previous results reported with SR-fMRI. Our results indicate that the sources of intermediate SR-fMRI FFA responses to nonfaces are neural populations within the FFA that are highly selective to nonfaces. Similarly, our data suggest that the source of adaptation for nonfaces¹³ in the FFA is the adaptation of specific neural populations that are selective to nonfaces rather than the adaptation of neural populations that respond to both faces and nonfaces. This proposition could be tested with future HR-fMRI experiments.

The main point illustrated by the models (**Fig. 1**) is that fundamentally different fine-scale functional organizations may yield identical SR-fMRI responses. Higher resolution imaging was necessary to distinguish between these functional organizations, as we demon-

strated in the present study. Our results suggest that caution should be exercised when making inferences about functional specialization from SR-fMRI data. First, SR-fMRI may average responses across heterogeneous neural populations. These effects are compounded by approaches that spatially blur the data and average across spatially normalized brains. Therefore, under these conditions, the average response from a region of interest may reflect the combined response of distinct neural populations, suggesting caution when interpreting the data. Second, a common practice in SR-fMRI is finding that the response to 'X' is greater than the response to 'Y' and interpreting this as a specialization for 'X'. However, our data demonstrate that this is only one of several interpretations—another equally plausible interpretation is that the region processes both 'X' and 'Y' but contains more neurons that process 'X'. Thus, it is necessary to examine both the fine-scale structure and the selectivity before making assertions about specialization.

Notably, the models for fine-scale structure (**Fig. 1**) are viable across cortex more generally. The 'patchy model' (**Fig. 1c**) could be a model for the functional organization for human MT (monkey electrophysiology predicts that motion-selective columns³⁷ are intermixed with disparity columns that are responsive to, but not selective for, motion direction³⁸). The 'highly heterogeneous model' (**Fig. 1d**) could be a model of orientation maps in human V1 (monkey electrophysiology predicts that vertical and horizontal orientations are over-represented compared to oblique orientations^{39,40}). Because the cortex contains fine structure that cannot be resolved with SR-fMRI, future HR-fMRI experiments are needed to reveal the fine-scale structure of the human cortex.

In sum, the results of the present study are important because they show, for the first time, the fine-scale structure of the FFA; they challenge the prevailing hypothesis that the FFA is a uniform brain area in which all neurons are face-selective; and they highlight the importance of conducting HR-fMRI to reveal the fine-scale structure of cortical activations in the human brain.

METHODS

Subjects. Five subjects (two female), ages 23–36, with corrected or normal vision, participated in these experiments. Subjects gave written informed consent to participate in the experiment. The experimental protocol was approved by Stanford University Institutional Review Board.

Scanning. Subjects were scanned on a research-only GE 3-Tesla Signa scanner at the Lucas Imaging Center at Stanford University. We imaged the right hemisphere using a custom-built surface coil (2.5" in diameter). The coil was worn on the lateral surface of the head, close to the brain regions that were imaged.

1. HR-fMRI. We acquired 12 1-mm-thick slices oriented in a sagittal-coronal plane, using a two-shot T2*-sensitive spiral acquisition sequence (FOV = 14 cm, TE = 25 ms, TR = 1s, flip angle = 62° and bandwidth = 125 kHz) to obtain 1 mm³ voxels. We also acquired T1-weighted in-plane anatomical images using a standard two-dimensional (2D) RF-spoiled GRASS (SPGR) sequence (minimum TE, TR = 100 ms, flip angle = 45°, bandwidth = 16 kHz) giving a voxel size of 0.5 mm × 0.5 mm × 1 mm. The in-plane anatomical images had the same prescription as the fMRI scans.

2. SR-fMRI. We conducted the same experiment on the same subjects using the same coil, but here we acquired data at a resolution of 3 mm × 3 mm × 3 mm. Scanning parameters were identical except for FOV = 16 cm.

3. Three-dimensional (3D) volume anatomicals. Subjects also participated in another scanning session in which a high-resolution anatomical volume of the whole brain was acquired with a head-coil using an inversion recovery SPGR sequence that yields good gray-white contrast. This anatomical volume was used for segmenting the brain into gray and white matter, and for visualization on the volume and inflated views.

4. Hierarchical pseudo-real-time scanning. To optimally position the HR-fMRI slices, subjects were first scanned with SR-fMRI during a functional localizer experiment. This experiment included alternating blocks (12 s each) of faces and scrambled faces. We used this scan to identify voxels that activated more strongly to faces than scrambled faces, using a correlation analysis. Prescriptions that included face-responsive voxels and also included either the fusiform gyrus or the OTS were chosen for subsequent high-resolution experiments.

Experimental design. The experiment included alternating blocks of gray-level images of faces, objects (four-legged animals, cars and abstract sculptures) and images randomly scrambled into 64 squares of 8×8 pixels (Supplementary Fig. 1). Each image subtended 15.2° of the subject's visual field, centered on the fovea. Blocks contained images of one category and were 12 s in duration; images appeared at a rate of 1 Hz. Animal images contained photographs of whole animals in which the face was typically in a nonfrontal view and occupied less than one-eighth of the image. In each nonscrambled block, 0–2 trials of scrambled images occurred randomly within the block. Subjects were asked to fixate and then, using a button box, to indicate the category of each image (face, animal, car, sculpture or scrambled). Each subject participated in eight runs (a run refers to a series of brain acquisitions in which we acquired responses to all experimental conditions) using different images and block orders to test the generality of our results. In total, subjects viewed 384 whole images and 400 scrambled images. Scrambled blocks were used as the baseline.

Data analysis. Data were analyzed in MATLAB (Mathworks) using the mrVISTA toolbox (<http://white.stanford.edu/software/>).

1. Anatomical data. Anatomical volumes were segmented into gray and white volumes. From this segmentation we reconstructed the cortical surface (gray matter) of each subject.

2. Functional data. Data were aligned to the high-resolution anatomical volume, which allowed us to compare regions of interest across scans (for example, $3 \text{ mm} \times 3 \text{ mm} \times 3 \text{ mm}$ and $1 \text{ mm} \times 1 \text{ mm} \times 1 \text{ mm}$ data) and to visualize activations on the volume as well as on the inflated brain. Functional data were motion-corrected and detrended. Standard general linear model (GLM) analyses were conducted to create voxel-by-voxel activation maps for various conditions (for example, Figs. 2 and 3). We did not spatially smooth the data at any resolution.

3. Estimation of the goodness of the alignment. To assess the goodness of our alignment (in-plane anatomicals to the volume anatomical) and correspondence across sessions, we chose, for each subject, two anatomical landmarks on in-plane anatomical images. These landmarks consisted of points of high curvature of the sulci or gyri in the vicinity of the fusiform gyrus. We projected these landmarks to the anatomical volume and measured their coordinates on the volume (x , y and z , in mm). We then calculated the average distance on the volume anatomical between the same landmark on $1 \text{ mm} \times 1 \text{ mm} \times 1 \text{ mm}$ and $3 \text{ mm} \times 3 \text{ mm} \times 3 \text{ mm}$ data. The average displacement was $1.6 \pm 1.1 \text{ mm}$. This analysis increased our confidence that data across scans were measured from the same cortical regions.

4. Definition of the FFA. The FFA was defined as a cluster of $3 \text{ mm} \times 3 \text{ mm} \times 3 \text{ mm}$ voxels in the fusiform gyrus that responded more strongly to faces than to inanimate objects with $P < 0.001$, at the voxel level. This is the standard way to define face-selective voxels in the human brain^{1–5}. The Talairach coordinates of the center of the FFA were $34 \pm 5 \text{ mm}$, $-52 \pm 8 \text{ mm}$ and $-13 \pm 4 \text{ mm}$ (mean \pm s.d. across subjects).

5. Amplitude calculation. For each voxel we calculated the amplitude of response as the average activation over stimulus blocks 6–12 s after trial onset minus the response during the baseline averaged over 6 s before each stimulus block. We chose these intervals as they do not include the rise and fall times of blood oxygenation level-dependent (BOLD) responses. Unless otherwise mentioned we used the scrambled condition as the baseline.

6. Selectivity of responses. For each voxel we determined the preferred category (that is, the category that yielded the maximum response) and then calculated its selectivity index for that preferred category:

$$\text{Selectivity index} = \frac{\text{Preferred} - \text{Nonpreferred}}{\text{Preferred} + |\text{Nonpreferred}|} \quad (1)$$

where 'Preferred' indicates the amplitude of the category that yielded the maximal response and 'Nonpreferred' indicates the average amplitude of other categories. A value of 1 indicates maximum selectivity, and a value of 0 indicates that there is no preference to any category. Therefore, the selectivity index codes what the preferred category is for a voxel, as well as how strong this preference is. From this calculation, we generated selectivity maps for FFA voxels (Fig. 2c, 4a and Supplementary Fig. 3). The hue reflects the preferred category and the saturation reflects the degree of selectivity (saturated colors indicate high selectivity, and faded colors indicate low selectivity).

7. Reliability analyses. We used several measures to examine whether responses recorded at high resolution are reliable. Each of these analyses was conducted for each subject.

(i) Reliability of activation patterns. For each subject and category (face, animal, car and sculpture), we created voxel-by-voxel activation maps of each category (category > scrambled, thresholded $P < 0.005$) separately for two independent data sets (odd and even runs) acquired in separate fMRI runs (that is, they did not include successive brain acquisitions) using different images. Then, for each subject and category, we calculated the cross-correlation between activations across the entire FFA for odd and even runs and tested whether this correlation was significant. During even and odd runs, subjects viewed different images, so that this calculation tests the reliability of responses for a category rather than reliability of responses to a particular set of images. The null correlation was estimated via a permutation-based method in which we randomly permuted the amplitudes of odd runs across voxels 60 times, and computed the average cross-correlation between permuted odd runs with original even runs. The null cross-correlation averaged across subjects and conditions was -0.002 ± 0.001 . We tested the significance of these correlations within and across subjects, as presented in the main text.

(ii) Reliability of time series. For each subject, we extracted the average time series separately for two data sets (odd and even runs) across the entire FFA and calculated the cross-correlation (and its significance) between the odd time series and the even time series. The time series for each condition (block and run) was extracted during a time window extending from 4 s before block onset to 22 s after block onset. We separately averaged responses across even-run blocks and odd-run blocks. Thus, the time series that was used to calculate the temporal reliability included 56 time points (4 conditions \times 14 time points per condition). We estimated the null temporal correlation by means of a permutation method in which we randomly permuted the time series for odd runs 60 times and computed the average cross-correlation between permuted odd runs with original even runs. The null cross-correlation was 0.0082 ± 0.0025 . We also tested the significance across subjects using a t -test.

(iii) Reliability of selectivity indices. We calculated selectivity indices (see equation (1)) for FFA voxels, separately for the odd runs and even runs. Then we calculated the cross-correlation between selectivity indices of FFA voxels for odd versus even runs and tested whether this correlation was significant both within and across subjects using t -tests.

(iv) Functional signal-to-noise ratio measurements. We calculated the functional signal-to-noise ratio as

$$\text{SNR} = 20 \log \left(\frac{E[\text{response}]}{\text{std}[\text{response}]} \right).$$

The numerator reflects the mean BOLD amplitude (versus baseline) and the denominator indicates the amplitude of the baseline fluctuations. We report functional signal-to-noise ratio for the amplitude of response to faces measured relative to the scrambled baseline for the entire region of interest (ROI) as defined with SR-fMRI.

8. Projection of HR-ROIs onto the inflated cortex (Fig. 5 and Supplementary Fig. 4). High-resolution data was acquired at 1 mm^3 , which is smaller than the cortical thickness. Therefore, unlike SR-fMRI voxels, which include all layers of the cortical sheet, voxels in HR-fMRI are small enough to project to various layers of the cortical sheet. To produce these figures, high-resolution voxels from all gray layers were projected onto a single surface using a nearest-neighbor metric. Color indicates the preferred category, but not the selectivity level.

9. PCA (ref. 17). PCA was conducted for each subject and resolution. For each subject, we calculated the amplitude of response in each FFA voxel, across

all runs for each category. This produces a matrix of four rows (each row is the pattern of activation across the FFA for one category) and n columns (each column is one voxel). We performed PCA of this data, by calculating the covariance matrix and performing singular value decomposition (SVD) on the covariance matrix. This analysis produces the eigenvectors (PCs), the variance associated with each component, and the scores (contribution of each component per voxel). The components are ordered according to their variance, starting from the component that accounts for the most variance in the data. In **Figure 6** we plotted the PCs that accounted for at least 95% of the variance of the data, for each subject and resolution.

10. Estimating SR-fMRI FFA response from HR-fMRI responses (**Fig. 7b**). We examined whether the maximal response to faces and intermediate response to nonfaces measured across the entire FFA can be explained by averaging over heterogeneous voxel populations. Therefore, we simulated the average FFA response measured with SR-fMRI as a weighted sum of selective high-resolution population responses (that is, face, animal, car and sculpture voxels). The response of each high-resolution population was weighted in proportion to the average number of voxels selective to that category:

$$\hat{\text{FFA}}(t) = A \frac{\sum_{i=1}^4 n_i h_i(t)}{\sum_{i=1}^4 n_i},$$

where $\hat{\text{FFA}}(t)$ is the estimated SR-fMRI FFA response; n_i is the average number of high-resolution voxels selective to category i ; and $h_i(t)$ is the average time series of the high-resolution population selective to category i . A is a constant.

Analysis of object-selective cortex. We used SR-fMRI to define object-selective regions using the standard definition for LOC (ref. 18; here we used cars and sculptures > scrambled, $P < 10^{-5}$, at the voxel level). Given the limited volume of our high-resolution prescription, only the anterior portion of the LOC (pFUS/OTS in previous studies^{16,19,20}; see **Fig. 8a**) was included in our high-resolution slices for 4 of 5 subjects. Data from this region were analyzed during standard-resolution and high-resolution scans using the same methods as described above for the FFA.

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We would like to thank B. Wandell, Y. Spector, N. Kanwisher and C. Baker for fruitful discussions; G. Golarai, D. Remus and D. Yoon for their comments on the manuscript; B. Dougherty for his help in developing software to project ROIs across gray layers; and J. Vinberg for his help in scanning subjects. This research was funded by grants from the National Eye Institute (1 R21EY016199-0) and the Whitehall Foundation (2005-05-111-RES) to K.G.-S. R.S. was supported by the National Eye Institute (5 F31 EY015937).

AUTHOR CONTRIBUTIONS

K.G.-S. contributed to all aspects of this experiment: design, data collection, development of data analysis methods, code development, and data analyses. R.S. contributed to the design, data collection, code and data analysis development, and data analyses. D.R. developed the surface coil used in these experiments and implemented and developed the HR-fMRI protocols used in these experiments. All authors contributed to understanding the implications of our results and to preparing the manuscript.

COMPETING INTERESTS STATEMENT

The authors declare that they have no competing financial interests.

Published online at <http://www.nature.com/natureneuroscience>

Reprints and permissions information is available online at <http://npg.nature.com/reprintsandpermissions/>

1. Kanwisher, N., McDermott, J. & Chun, M.M. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J. Neurosci.* **17**, 4302–4311 (1997).
2. George, N. *et al.* Contrast polarity and face recognition in the human fusiform gyrus. *Nat. Neurosci.* **2**, 574–580 (1999).
3. Tong, F., Nakayama, K., Vaughan, J.T. & Kanwisher, N. Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron* **21**, 753–759 (1998).

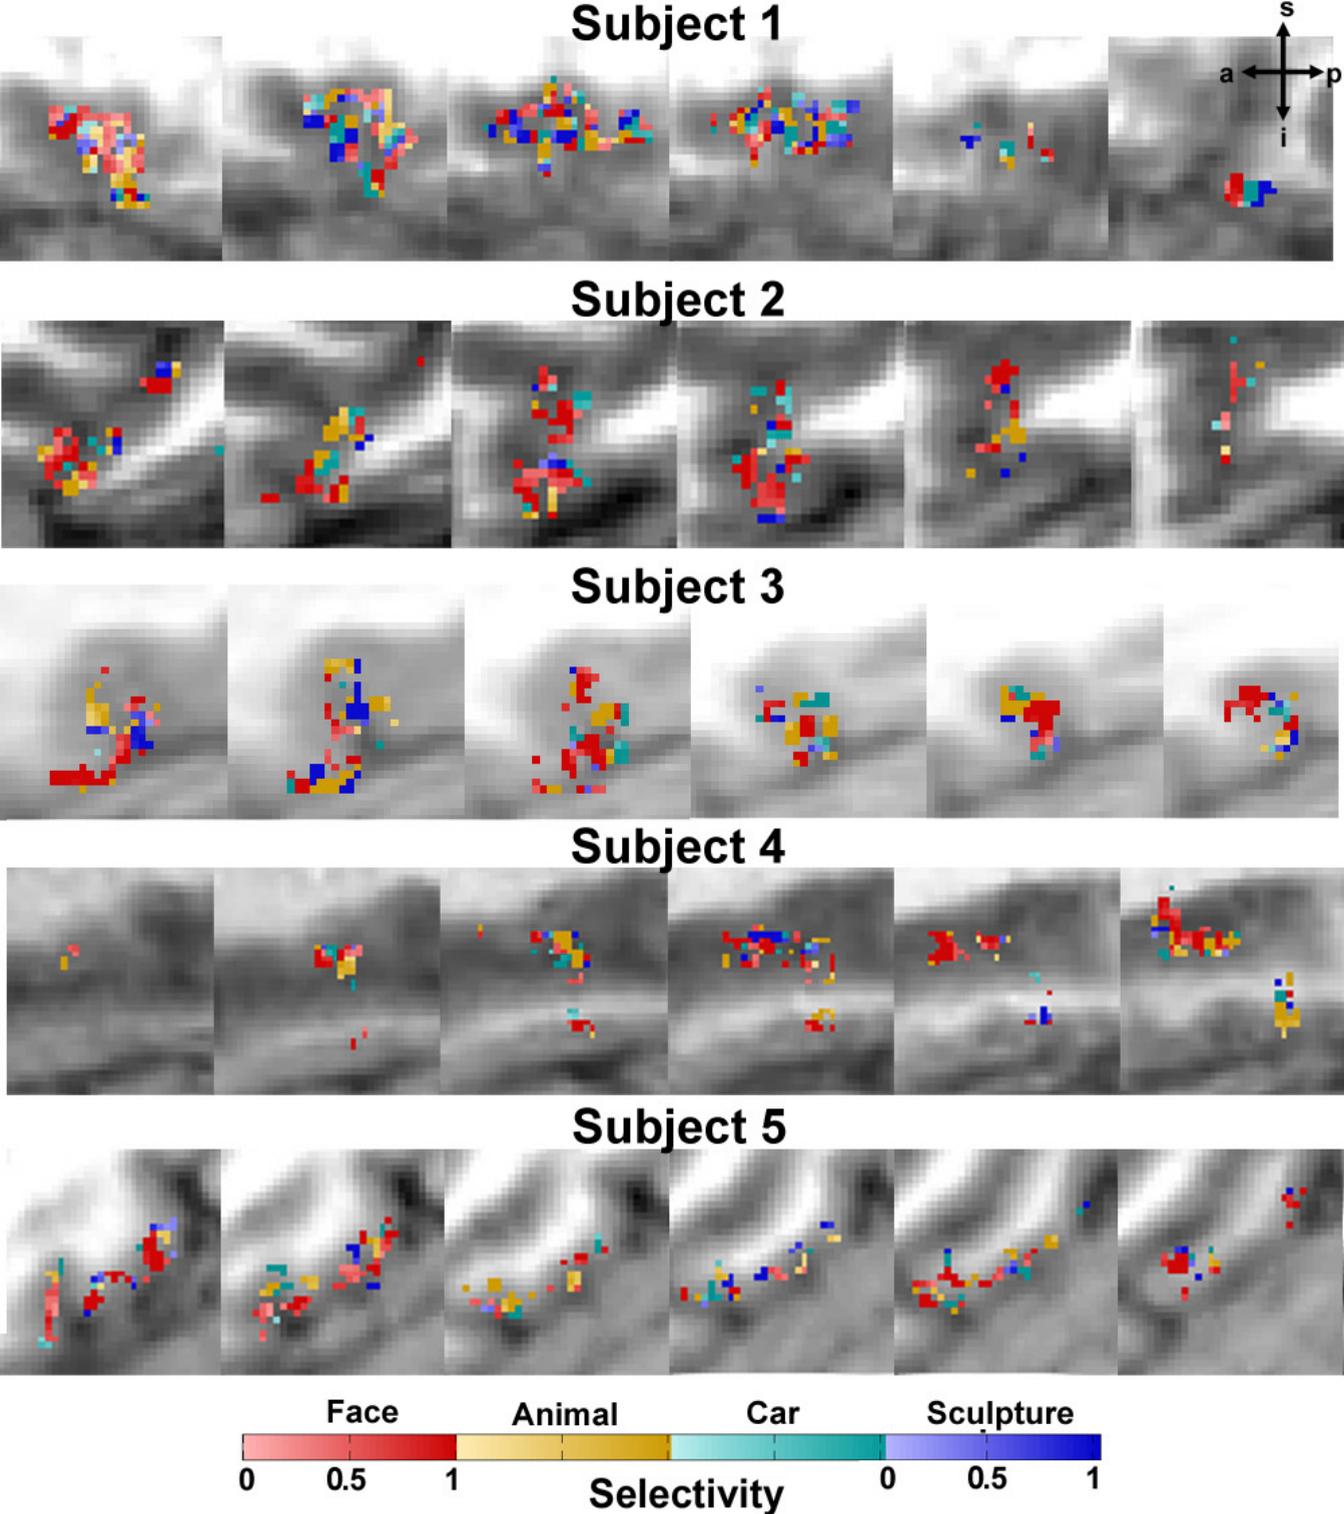
4. Hasson, U., Hendler, T., Ben Bashat, D. & Malach, R. Vase or face? A neural correlate of shape-selective grouping processes in the human brain. *J. Cogn. Neurosci.* **13**, 744–753 (2001).
5. Grill-Spector, K., Knouf, N. & Kanwisher, N. The fusiform face area subserves face perception, not generic within-category identification. *Nat. Neurosci.* **7**, 555–562 (2004).
6. Kanwisher, N. Domain specificity in face perception. *Nat. Neurosci.* **3**, 759–763 (2000).
7. Schwarzlose, R.F., Baker, C.I. & Kanwisher, N.K. Separate face and body selectivity on the fusiform gyrus. *J. Neurosci.* **25**, 11055–11059 (2005).
8. Winston, J.S., Henson, R.N., Fine-Goulden, M.R. & Dolan, R.J. fMRI-adaptation reveals dissociable neural representations of identity and expression in face perception. *J. Neurophysiol.* **92**, 1830–1839 (2004).
9. Spiridon, M. & Kanwisher, N. How distributed is visual category information in human occipito-temporal cortex? An fMRI study. *Neuron* **35**, 1157–1165 (2002).
10. Edelman, S., Grill-Spector, K., Kushnir, T. & Malach, R. Toward direct visualization of the internal shape representation space by fMRI. *Psychobiology* **26**, 309–321 (1998).
11. Haxby, J.V. *et al.* Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* **293**, 2425–2430 (2001).
12. Ishai, A., Ungerleider, L.G., Martin, A., Schouten, J.L. & Haxby, J.V. Distributed representation of objects in the human ventral visual pathway. *Proc. Natl. Acad. Sci. USA* **96**, 9379–9384 (1999).
13. Avidan, G., Hasson, U., Hendler, T., Zohary, E. & Malach, R. Analysis of the neuronal selectivity underlying low fMRI signals. *Curr. Biol.* **12**, 964–972 (2002).
14. Desimone, R., Albright, T.D., Gross, C.G. & Bruce, C. Stimulus-selective properties of inferior temporal neurons in the macaque. *J. Neurosci.* **4**, 2051–2062 (1984).
15. Tsao, D.Y., Freiwald, W.A., Tootell, R.B. & Livingstone, M.S. A cortical region consisting entirely of face-selective cells. *Science* **311**, 670–674 (2006).
16. Grill-Spector, K., Kourtzi, Z. & Kanwisher, N. The lateral occipital complex and its role in object recognition. *Vision Res.* **41**, 1409–1422 (2001).
17. Duda, R.O., Hart, P.E. & Stork, D.G. *Pattern Classification* (John Wiley & Sons, New York, 2001).
18. Malach, R. *et al.* Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proc. Natl. Acad. Sci. USA* **92**, 8135–8139 (1995).
19. Grill-Spector, K. The neural basis of object perception. *Curr. Opin. Neurobiol.* **13**, 159–166 (2003).
20. Sayres, R. & Grill-Spector, K. Object-selective cortex exhibits performance-independent repetition suppression. *J. Neurophysiol.* **95**, 995–1007 (2006).
21. Peelen, M.V. & Downing, P.E. Selectivity for the human body in the fusiform gyrus. *J. Neurophysiol.* **93**, 603–608 (2005).
22. Saleem, K.S. *et al.* Magnetic resonance imaging of neuronal connections in the macaque monkey. *Neuron* **34**, 685–700 (2002).
23. Tolias, A.S. *et al.* Mapping cortical activity elicited with electrical microstimulation using fMRI in the macaque. *Neuron* **48**, 901–911 (2005).
24. Tsao, D.Y., Freiwald, W.A., Knutsen, T.A., Mandeville, J.B. & Tootell, R.B. Faces and objects in macaque cerebral cortex. *Nat. Neurosci.* **6**, 989–995 (2003).
25. Pinsk, M.A., DeSimone, K., Moore, T., Gross, C.G. & Kastner, S. Representations of faces and body parts in macaque temporal cortex: a functional MRI study. *Proc. Natl. Acad. Sci. USA* **102**, 6996–7001 (2005).
26. Gauthier, I., Tarr, M.J., Anderson, A.W., Skudlarski, P. & Gore, J.C. Activation of the middle fusiform ‘face area’ increases with expertise in recognizing novel objects. *Nat. Neurosci.* **2**, 568–573 (1999).
27. Tarr, M.J. & Gauthier, I. FFA: a flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nat. Neurosci.* **3**, 764–769 (2000).
28. Levy, I., Hasson, U., Avidan, G., Hendler, T. & Malach, R. Center-periphery organization of human object areas. *Nat. Neurosci.* **4**, 533–539 (2001).
29. Malach, R., Levy, I. & Hasson, U. The topography of high-order human object areas. *Trends Cogn. Sci.* **6**, 176–184 (2002).
30. Wang, G., Tanaka, K. & Tanifuji, M. Optical imaging of functional organization in the monkey inferotemporal cortex. *Science* **272**, 1665–1668 (1996).
31. Grill-Spector, K. *et al.* Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron* **24**, 187–203 (1999).
32. Tsunoda, K., Yamane, Y., Nishizaki, M. & Tanifuji, M. Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns. *Nat. Neurosci.* **4**, 832–838 (2001).
33. Gauthier, I. & Nelson, C.A. The development of face expertise. *Curr. Opin. Neurobiol.* **11**, 219–224 (2001).
34. Golby, A.J., Gabrieli, J.D., Chiao, J.Y. & Eberhardt, J.L. Differential responses in the fusiform region to same-race and other-race faces. *Nat. Neurosci.* **4**, 845–850 (2001).
35. Fujita, I., Tanaka, K., Ito, M. & Cheng, K. Columns for visual features of objects in monkey inferotemporal cortex. *Nature* **360**, 343–346 (1992).
36. Olshausen, B.A. & Field, D.J. Sparse coding of sensory inputs. *Curr. Opin. Neurobiol.* **14**, 481–487 (2004).
37. Albright, T.D., Desimone, R. & Gross, C.G. Columnar organization of directionally selective cells in visual area MT of the macaque. *J. Neurophysiol.* **51**, 16–31 (1984).
38. DeAngelis, G.C. & Newsome, W.T. Organization of disparity-selective neurons in macaque area MT. *J. Neurosci.* **19**, 1398–1415 (1999).
39. Coppola, D.M., White, L.E., Fitzpatrick, D. & Purves, D. Unequal representation of cardinal and oblique contours in ferret visual cortex. *Proc. Natl. Acad. Sci. USA* **95**, 2621–2623 (1998).
40. Chapman, B. & Bonhoeffer, T. Overrepresentation of horizontal and vertical orientation preferences in developing ferret area 17. *Proc. Natl. Acad. Sci. USA* **95**, 2609–2614 (1998).





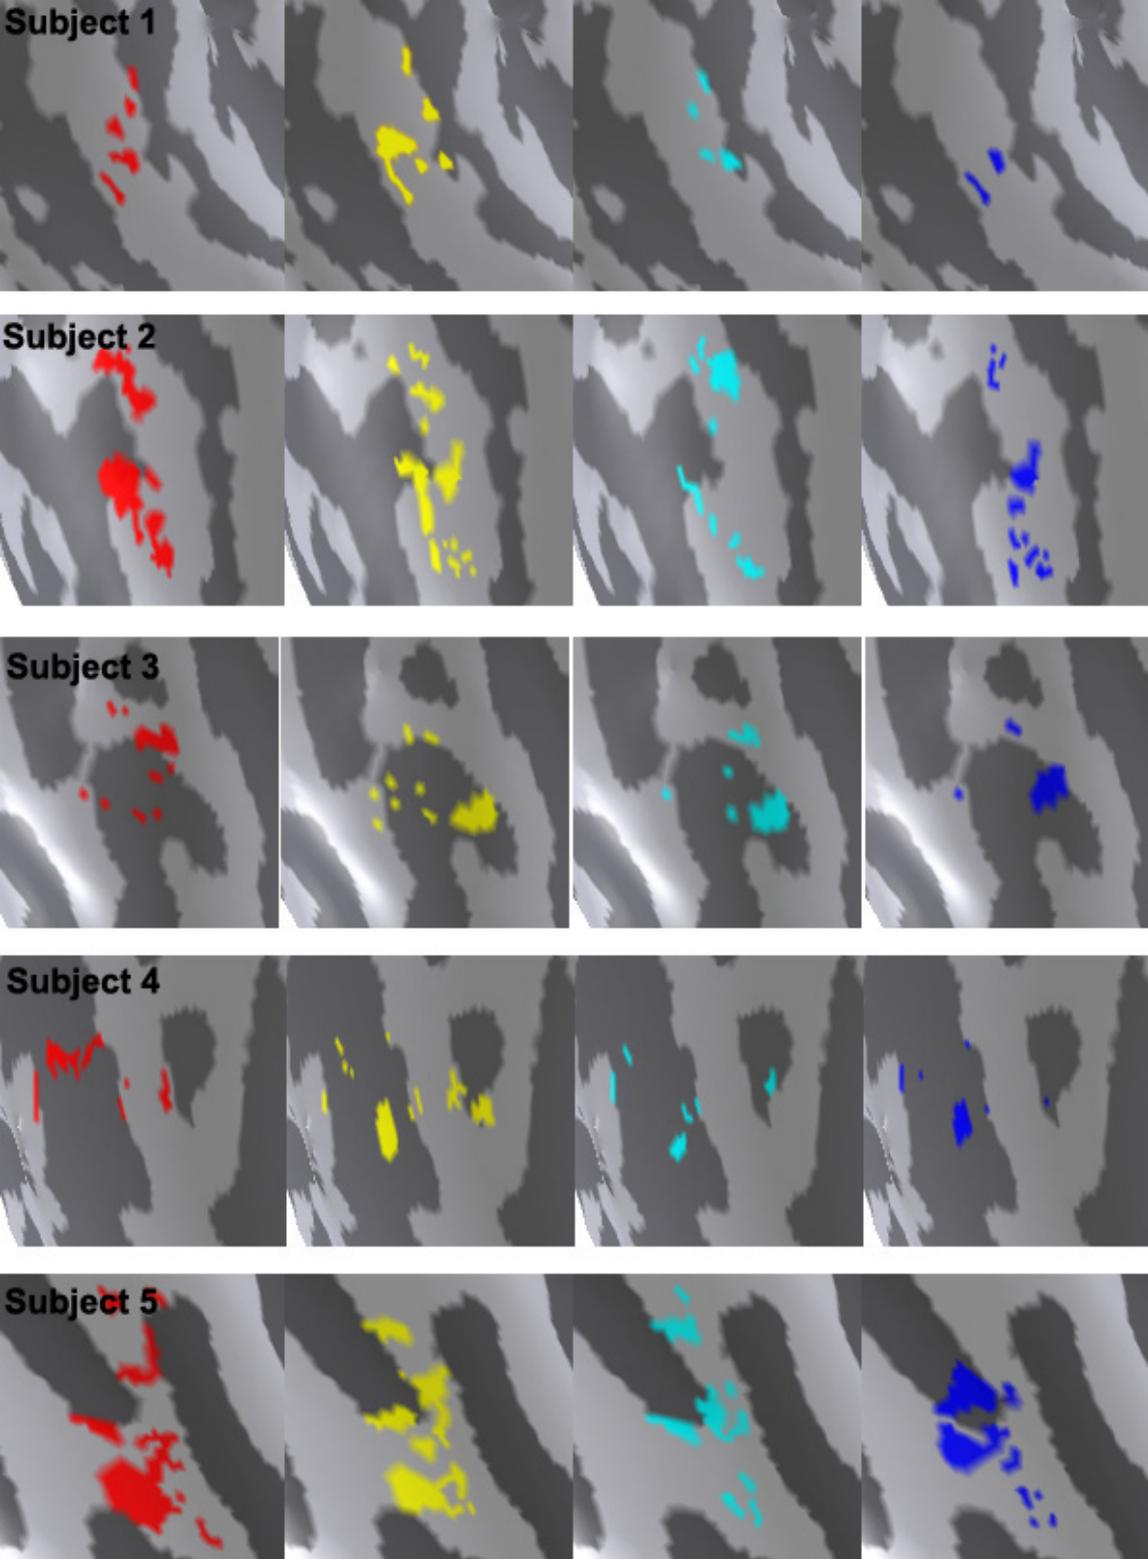
Supplementary Figure S1: Examples of stimuli used in the experiment.

From top to bottom: faces, animals, cars, sculptures and scrambled. Red cross indicates fixation.



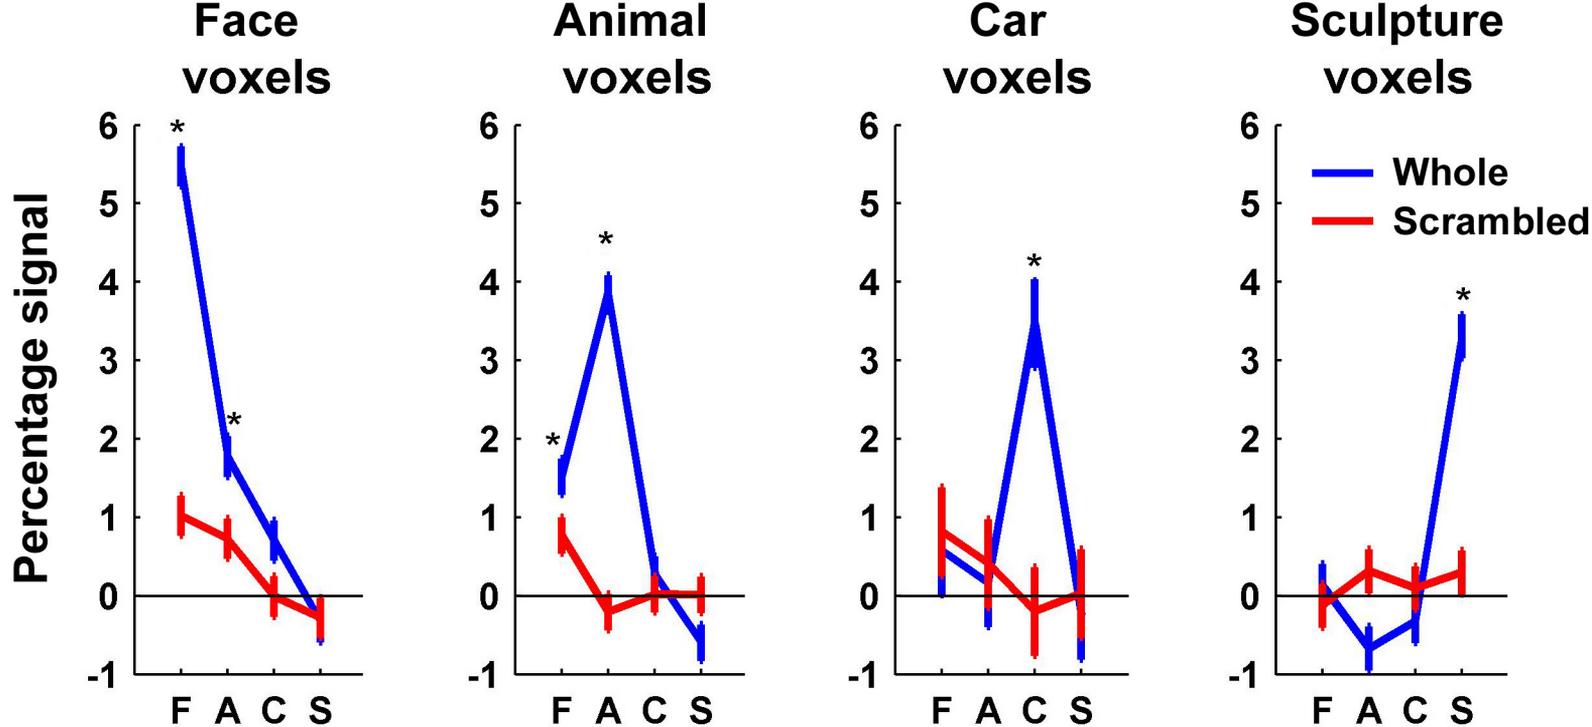
Supplementary Figure 3: Individual subject selectivity maps

Each row contains consecutive inplanes from a single subject, zoomed on the FFA. Slices were prescribed in an oblique sagittal direction (see Methods). Slices are arranged from lateral to medial (left -> right). Colors indicate preferred category. Hue indicates level of selectivity, as indicated by the colorbar. Acronyms: a: anterior, p: posterior, i: inferior, s: superior.



Supplementary Figure 4: Single subject category selective subregions on the inflated cortical surface.

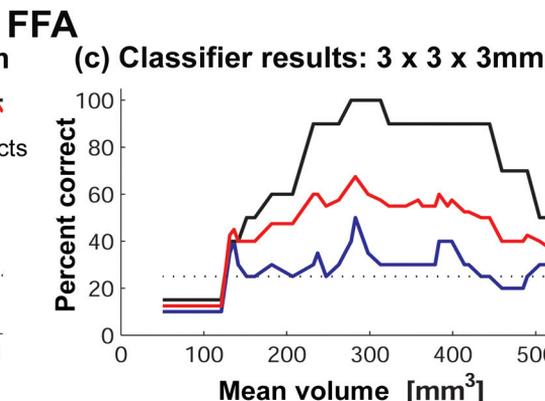
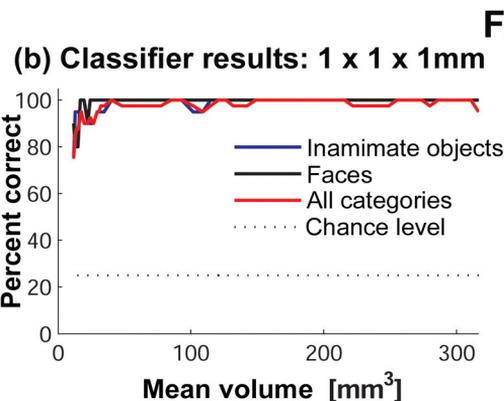
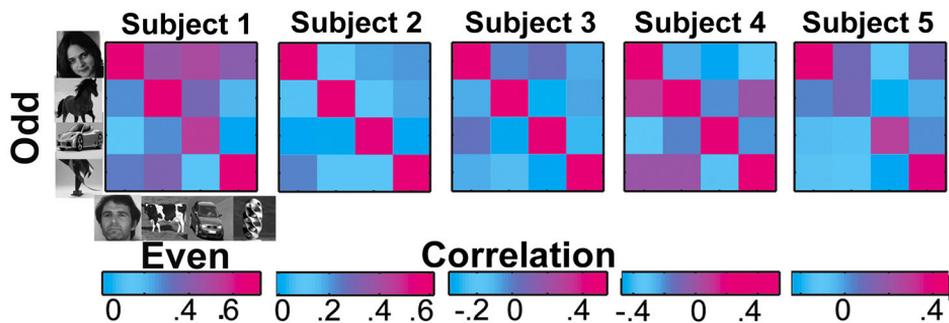
Each row contains single subject maps, zoomed around the fusiform gyrus and OTS. Colored regions represent category selective subregions (HR voxels just preferred category). Red: face-selective voxels; Yellow: animal-selective; Cyan: car-selective; Blue: sculpture-selective. Data from all layers were projected to the cortical surface. Some patches selective to different categories may appear overlapping on the cortical surface, although they are not overlapping in the inplane slices or volume anatomy (see Supplementary Fig. 3). The apparent overlap on the inflated cortical surface occurs because: (i) HR voxels are smaller than cortical thickness and (ii) there is a one-to-many mapping between the anatomical volume and the inflated cortical surface.



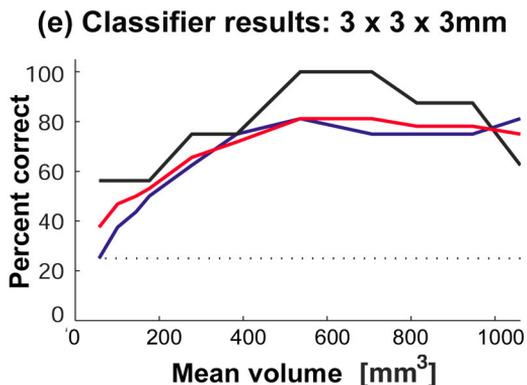
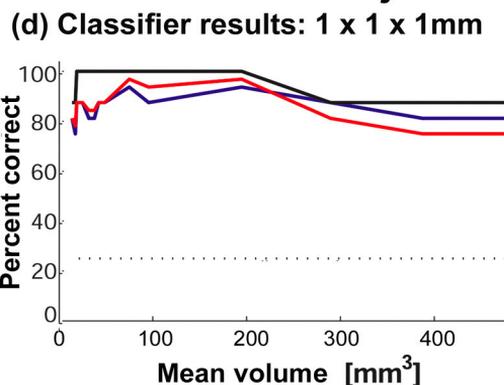
Supplementary Figure 6: FFA responses to whole and scrambled images versus fixation.

Data were acquired at 1x1x1mm and are averaged across three subjects. Blue: Average response amplitude to whole images of a category versus the fixation baseline. Red: Average response amplitude to scrambled images of a category versus the fixation baseline. Error bars: SEM across subjects. X-axis: stimulus category. F: faces; A: animals; C: cars; S: sculptures. Asterisks: significantly larger than baseline, t-test across subjects, $P < 0.05$;

(a) FFA Pattern analysis: 1 x 1 x 1mm data



Object-Selective Cortex



Supplementary Figure 5: Decoding distributed FFA activation patterns.

(a) Crosscorrelation between distributed HR-voxel amplitudes across the FFA for odd vs. even runs. Each panel shows data from one subject. (b-e) Results of a winner-take-all (WTA) classifier. (b-c) FFA data. WTA classifier hit rate as a function of FFA volume, for faces (black), inanimate objects (blue) and across all categories (red). Data are averaged across 5 subjects. Voxels are ranked in descending order according to the reliability of their responses (details in Supplementary Notes). (d-e) Object-selective cortex (from the anterior part of LOC overlapping the OTS). Data are averaged across 4 subjects. Conventions same as (b-c).

Supplementary Table 1: Reliability Analyses of HR fMRI.

Results for the entire region of interest as defined with SR-fMRI. In Exp.1 baseline was scrambled. In Exp. 2 baseline was a blank screen with fixation.

	SNR Face > baseline	Time course odd vs. even	Distributed Activation odd vs. even	Selectivity map odd vs. even
Exp 1. Object-selective (OTS)	15 ± 2 dB	$\rho = 0.56 \pm 0.23$ $P < 0.001$	$\rho_{\text{face}} = 0.50 \pm 0.09$ $\rho_{\text{animal}} = 0.31 \pm 0.10$ $\rho_{\text{car}} = 0.40 \pm 0.07$ $\rho_{\text{sculpture}} = 0.38 \pm 0.05$ $P < 0.01$	$\rho = 0.84 \pm 0.1$ $P < 0.00001$
Exp. 2 FFA	18 ± 2 dB	$\rho = 0.57 \pm 0.13$ $P < 0.001$	$\rho_{\text{face}} = 0.65 \pm 0.02$ $\rho_{\text{animal}} = 0.56 \pm 0.05$ $\rho_{\text{car}} = 0.52 \pm 0.08$ $\rho_{\text{sculpture}} = 0.42 \pm 0.10$ $P < 0.00001$	$\rho = 0.94 \pm 0.01$ $P < 0.00001$

SUPPLEMENTARY NOTE

Pattern Analyses

We performed a pattern analysis that examines both the reliability and distinctiveness of distributed activation patterns. We calculated for each FFA voxel its response amplitude for each category, separately for odd and even scans. This calculation produces a matrix whose rows consist of the distributed pattern of response for different categories. Then we calculated the crosscorrelation matrix of odd vs. even distributed responses.

Pattern analysis of HR-fMRI revealed that the distributed activations across the FFA are both reliable (high on-diagonal values) and distinct (low off-diagonal values, **Supplemental Fig. 5a**). This was evident in each subject and for all categories. Therefore, we found reliable patterns of activation across the FFA for both faces and nonfaces, and different categories elicited distinct activations across the FFA.

Decoding Distributed FFA Activation Patterns

Methods

Winner Take-All (WTA) Classifier: We implemented a WTA classifier to examine whether an independent classifier can determine from the distributed activation patterns what category subjects viewed. (The distributed activation for a category is a vector of response amplitudes of each voxel to that particular category.) The WTA is trained with the distributed response pattern for each category from one data set (e.g. odd runs in one subject). During classification, it receives as input the distributed amplitudes from another data set (e.g. even runs in the same subject). It calculates the cross correlation between the input and each of the training pattern sets. It determines the category based on the comparison that yielded highest correlation between the input pattern and the training set. This analysis was repeated twice, once training with odd runs and testing with even runs and once training with even and testing with odd runs (cross-correlation

matrix is not symmetric). Results are averaged across these two sets (Supplementary Fig. S5). Chance level is 25%.

Applying WTA Classifier to Progressive subsets of voxels: In order to assess the robustness of the classifier across voxels and subjects, we ran the WTA Classifier on progressive subsets of the data (**Supplementary Fig. 5**). Voxels were ranked in descending order according to the reliability of each voxel (the crosscorrelation between 4-dimensional vectors of response amplitudes from even and odd runs). This is an independent measure to the distributed activation pattern.

Results

FFA

The WTA classifier was able to determine from the distributed activation across HR FFA voxels what category subjects viewed with 100% success (**Supplementary Fig. 5b**). This was true for both faces and nonfaces. The classifier was at 100% performance when using information from many ranges of subpopulations of HR-voxels in the FFA. Surprisingly, with only 22 HR-voxels (volume < one standard voxel) the classifier's performance was 92%. These data indicate that the category can be inferred from the response of a small number of HR-voxels within the FFA. PCA (**Fig. 6**) confirmed our pattern analyses and suggests that additional linear classifiers based on HR data will be successful.

With SR-fMRI, the WTA classifier was successful at identifying face stimuli from FFA data, but failed in discriminating between inanimate objects (Supplementary Fig. S5b), replicating previous results⁸. The classifier failed at standard resolution, because SR-voxels were uniformly selective for faces, and there was insufficient systematic variation to distinguish different nonface categories (see Fig. 6b). If the patterns to different categories do not vary systematically, linear classifiers cannot separate multiple categories.

Results of the classifier approach confirm our analyses presented in the main text that show that distributed HR-fMRI activation across the FFA (for different categories) are reliable and distinct. Furthermore, we show that an external classifier can decode from the patterns of HR-FFA activation what category subjects viewed, indicating that the distributed HR FFA activation contains information about both faces and nonfaces.

Object-Selective Cortex.

We applied PCA and the WTA classifier data from object-selective cortex (see Methods) obtained with SR and HR-fMRI.

SR-fMRI of object-selective cortex show that there is weak, but heterogeneous selectivity to different categories in individual voxels' responses. Indeed PCA of object-selective cortex shows that three PCs were necessary for explaining the variance in this data (% variance accounted by PCs: 67 ± 11 , 89 ± 4 , 11 ± 4). WTA applied to SR-fMRI data indicates that these biases are reliable and there is a range of voxels in which the WTA's performance is above chance.

Similar to SR-fMRI three PCs explained the variance of HR-fMRI of object-selective cortex (% variance accounted by PCs: 48 ± 7 , 28 ± 4 , 24 ± 4). However, PCs of HR-fMRI were more distinct than SR-fMRI. Therefore, WTA classifier performance was improved for HR object-selective cortex data. Similar to FFA data, the WTA performance was high for many subsets of HR voxels.

Overall, these pattern analyses show that distributed activations to specific categories are reliable, activations to different categories are distinct and there is information about both faces and objects in each of the FFA and object-selective cortex. Moreover, these trends are more clearly revealed by HR-fMRI than SR-fMRI.

Role of Local Features in Explaining Category-Selective Responses

Three of our subjects (subjects 2, 3 & 4) participated in a second experiment in which we examined whether category-selective clusters revealed at high resolution responded to global structure of images of a category or local features (that may be correlated to specific categories).

Methods

In the second experiment, subjects viewed in separate 12s blocks whole or scrambled images of one category which alternated with 12s blocks in which a fixation point appeared on a gray screen. Subjects were instructed to fixate. During each image block 0-2 randomly selected images were repeated and subjects were requested to press a button when an image repeated, During fixation blocks subjects pressed a button when the fixation changed color. Subjects participated in eight such runs that included different images and block order.

Data were motion corrected, detrended, and aligned to the volume anatomy of each subject. For each subject we projected the FFA (as defined from SR-fMRI data in the first experiment) onto inplane data. Results of the second experiment were reliable (**Supplementary Table 1: Exp 2 FFA**).

Results

Selectivity was defined similarly to Experiment 1 based on responses to whole images only (Equation 1) with the exception that percent signal was calculated relative to the fixation baseline. We found a similar profile of selectivity in Experiments 1 & 2 (compare **Supplemental Fig. 6-blue and Fig. 4b**). Similar to the first experiment we found more face-selective HR voxels than voxels selective to other categories in the FFA: on average there were 175 ± 21 face HR-voxels, 98 ± 30 animal, 68 ± 17 car and 63 ± 20 sculpture HR-voxels. Therefore, the selectivity of the response patterns in Experiment 2

is consistent with Experiment 1, even when activations were compared to a fixation baseline.

We then extracted responses from category-selective patches during scrambled blocks (**Supplemental Fig. 6**-red). In all category-selective patches we found a strong scrambling effect, in which responses to whole images of the preferred category were significantly higher than each of the scrambled categories (t-test, $P < 0.00001$ within-subject; $P < 0.05$ across-subject). There was an overall weak trend for higher responses during scrambled face blocks than other scrambled stimuli, but the difference between responses to different scrambled categories did not reach statistical significance in any of the category-selective clusters ($F_{1,3} < 1.96$, $P > 0.19$). We found a small, but significantly higher response for scrambled faces than scrambled images of other categories (t-test, $P < 0.05$) for 2 of the 3 subjects, but not for any other subjects, categories or clusters. Finally, a two-way analysis of variance using factors of image type (whole/scrambled) and category (face/animal/car/sculpture) revealed a significant interaction between image type and category ($F_{3,23} > 10$, $P < 0.0005$), suggesting a different pattern of selectivity for whole and scrambled images. Given the strong interaction between image scrambling and category and the significantly lower responses to all scrambled images, these data suggest that selectivity to specific categories is not driven by local features that are correlated with these categories.