

Category Selectivity in the Ventral Visual Pathway Confers Robustness to Clutter and Diverted Attention

Leila Reddy^{1,2,*} and Nancy Kanwisher^{1,*}

¹McGovern Institute for Brain Research
Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, Massachusetts 02139

Summary

Are objects coded by a small number of neurons or cortical regions that respond preferentially to the object in question, or by more distributed patterns of responses, including neurons or regions that respond only weakly? Distributed codes can represent a larger number of alternative items than sparse codes [1–3] but produce ambiguities when multiple items are represented simultaneously (the “superposition” problem) [1, 4]. Recent studies found category information in the distributed pattern of response across the ventral visual pathway, including in regions that do not “prefer” the object in question [5–8]. However, these studies measured neural responses to isolated objects, a situation atypical of real-world vision, where multiple objects are usually present simultaneously (“clutter”). We report that information in the spatial pattern of fMRI response about standard object categories is severely disrupted by clutter and eliminated when attention is diverted. However, information about preferred categories in category-specific regions is undiminished by clutter and partly preserved under diverted attention. These findings indicate that in natural conditions, the pattern of fMRI response provides robust category information only for objects coded in selective cortical regions and highlight the vulnerability of distributed representations to clutter [1, 2] and the advantages of sparse cortical codes in mitigating clutter costs.

Results

Ten subjects viewed blocks of objects from a given category presented either in isolation or in the presence of another simultaneously presented object (Figure 1). In the latter case, the object in question was either attended or unattended. Two categories that selectively activate specific extrastriate regions (faces and houses in the fusiform face area [FFA] [9] and parahippocampal place area [PPA] [10], respectively), and two categories that produce no strongly selective responses detectable with current fMRI resolution (shoes and cars) were used

[11]. To decode category information, pattern classification analyses [12, 13], including “correlation” analyses [5] and support vector machines [14] (Figure S6 in the Supplemental Data available online), were applied to three regions of interest: the FFA, PPA, and “ORX” (see Experimental Procedures). Subjects’ behavioral performance and eye-movement data is shown in Figures S1, S4, and S5.

fMRI Classification Performance for Isolated Stimuli

The category of an isolated object can be determined from the spatial profile of fMRI response [5]. Accordingly, in the isolated condition (Figure 2, orange bars), all four stimulus categories could be discriminated from each other above chance based on the pattern of response in each ROI, except cars in the FFA and PPA. Thus, we found above-chance discrimination of a nonpreferred category (shoes) in the FFA and the PPA, an effect compatible with Haxby et al. [5], but not found by [6] and weak in [8]. This discrimination for non-preferred stimuli was not driven by faces and houses; performance on discriminating only shoes from cars was still above chance (FFA: 70% [$t(9) = 2.57$; $p < .05$]; PPA: 81% [$t(9) = 3.59$; $p < .01$]).

Discrimination performance in the isolated condition depended on stimulus category and ROI: an ANOVA on classification performance over these two factors revealed a main effect of category ($F(3,27) = 17.26$; $p < .0001$). Following up on this effect, performance for “special” categories (faces and houses) was higher than for “standard” categories (cars and shoes) in all three regions (standard versus special category, FFA: $F(1,18) = 31.6$; $p < .0001$; PPA: $F(1,18) = 63.7$; $p < .0001$; ORX: $F(1,18) = 4.7$; $p < 0.05$). Further, there was a significant interaction between stimulus category and ROI ($F(6,54) = 3.9$; $p < 0.005$), indicating that the advantage for “special” categories was greater in the FFA and PPA than in ORX. Note that these results do not simply reflect variations in the global BOLD activity across each ROI (Figure S2).

Therefore, these data replicate previous results showing high classification performance based on the spatial pattern of fMRI response, and higher performance for “special” than “standard” categories, while further showing some discriminative information for non-preferred stimuli in the FFA and PPA.

Effects of Clutter

Is the discrimination performance observed in each ROI robust to clutter—i.e., when another unattended object is present simultaneously (“attended” condition)? A three-way ANOVA of ROI (FFA/PPA/ORX) \times category (faces/houses/shoes/cars) \times isolated/attended showed a main effect of significantly lower performance for the attended compared to the isolated condition (“clutter cost”) ($F(1,9) = 26.1$; $p < .001$) and a significant main effect of category ($F(3,27) = 63.9$; $p < .0001$), reflecting the

*Correspondence: lreddy@klab.caltech.edu (L.R.), ngk@mit.edu (N.K.)

²Present address: Université de Toulouse, CerCo, UPS, 31062 Toulouse, France, and CNRS, UMR549, Faculté de Médecine de Rangueil, 31062 Toulouse, France.

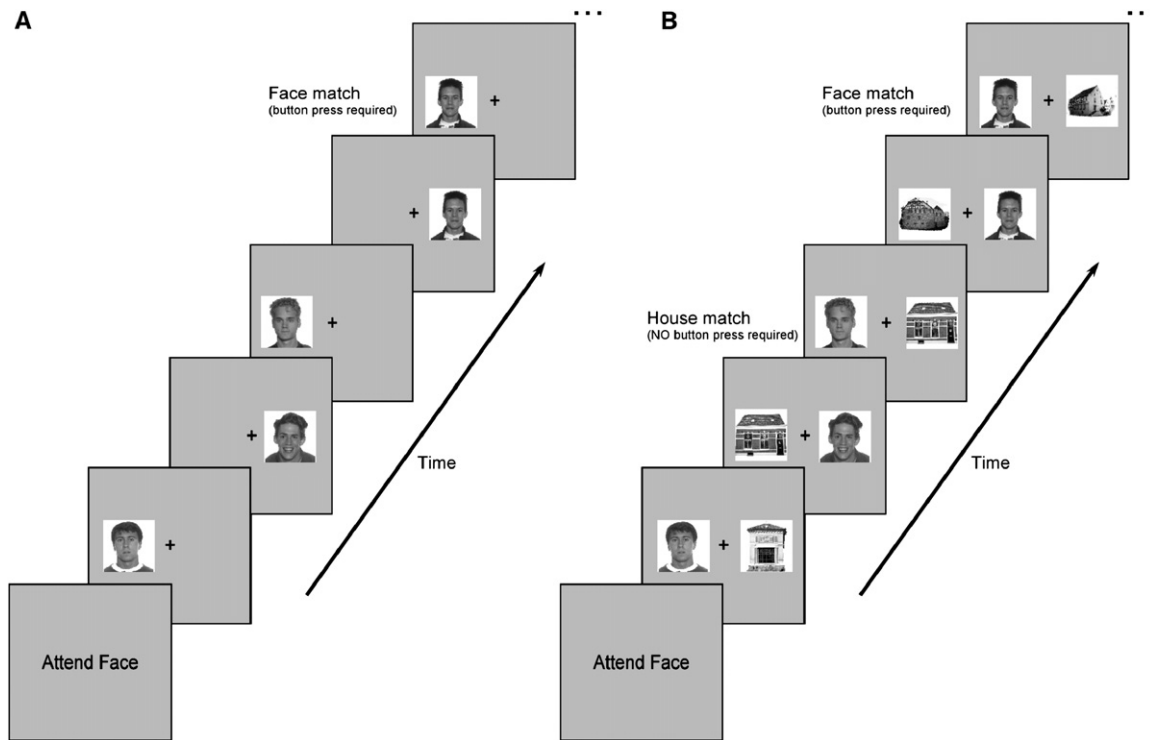


Figure 1. Experimental Design

Subjects were presented with either one ([A]; isolated condition) or two ([B]; attended and unattended conditions) streams of images that alternated on either side of fixation. Each block began with an instruction screen presented for 2 s telling subjects which category of stimuli they were to perform a 1-back task on. Images from four categories (faces, houses, shoes, cars) were used during the experiment. In the attended and unattended conditions, all possible pairs of image categories were presented to subjects.

previously described higher performance for faces and houses versus shoes and cars (Figure 2).

The effects of stimulus category and presentation condition (isolated versus attended) depended on ROI, as revealed by a double interaction of ROI × category

($F(6,54) = 12.23$; $p < .0001$) and a triple interaction of ROI × category × presentation condition ($F(6,54) = 4.52$; $p < 0.001$). To investigate classification performance within each ROI, a 2-way ANOVA of category by condition (isolated versus attended) was computed. The FFA




Figure 2. Effect of Clutter and Attention

fMRI discrimination performance and standard error across 10 subjects for discriminations involving faces, houses, shoes, and cars in the isolated (orange), attended (green), and unattended (yellow) stimulus-presentation conditions. The performance in the attended and unattended conditions reflects the average performance over pairs of each object category presented with all other objects. These results are shown for FFA (left), PPA (middle), and object responsive voxels with face-selective and house-selective voxels excluded (ORX) (right). Chance is at 50%. Error bars correspond to the standard error of the mean (SEM).

showed a main effect of category ($F(3,27) = 38.69$; $p < .0001$) and condition ($F(1,9) = 53.2$; $p < .0001$) and a significant interaction ($F(3,27) = 3.28$; $p < .05$). The PPA showed a main effect of category ($F(3,27) = 30.84$; $p < .0001$) and condition ($F(1,9) = 5.21$; $p < .05$), although the interaction was not significant ($F(3,27) = 2.03$; $p = .13$). The ORX showed a consistent clutter cost ($F(1,9) = 42.2$; $p < .0001$) across all categories ($F(3,27) = 17.1$; $p < .0001$), with no interaction of category by isolated/attended ($F(3,27) = .5$; $p > .5$).

The interaction effect in the FFA indicates that the clutter cost depends on object category. Pair-wise comparisons for faces and houses in the FFA and PPA showed that classification performance for non-preferred stimuli dropped substantially for the attended versus the isolated condition ($t(9) = 6.67$; $p < .0001$ for houses in the FFA and $t(9) = 4.37$; $p < .01$ for faces in the PPA). However, there was no significant drop for preferred stimuli in these regions ($t(9) = 0.26$; $p > .5$ for houses in the PPA and $t(9) = 1.88$; $p > .05$ for faces in the FFA). In contrast, in ORX classification performance dropped significantly for the attended compared to the isolated condition for all stimulus categories (faces: $t(9) = 4.23$; $p < .01$; houses: $t(9) = 4.06$; $p < .01$; shoes: $t(9) = 2.42$; $p < .05$; cars: $t(9) = 2.72$; $p < .05$), although it remained above chance. Note that when the “attended” condition was used as a decoding reference instead of the “isolated” condition, we observed the same pattern of results (Figure S7).

To directly test the significance of the greater clutter cost for faces and houses in ORX than in the category-selective regions, we performed a new 3-way ANOVA of condition (isolated/attended), ROI (the region that prefers the category in question versus ORX), and category (F versus H). A significant interaction of ROI \times condition ($F(1,9) = 14.02$; $p < .005$) indicated significantly greater sparing of face and house classification from clutter in their respective category-selective regions, versus ORX. A nonsignificant triple interaction ($F(1,9) = 3.01$; $p > .05$), indicated that this clutter sparing in category-selective regions did not differ for faces in the FFA versus houses in the PPA. Thus, category-selective regions selectively “protect” their preferred stimuli from the clutter cost that is devastating to other stimulus categories in each ROI. Similar results were obtained when classification performance in the isolated condition was equalized across all categories (Figure S3), and when we used a ROI-free approach (Figure S8).

Effects of Attention

The effects of attention were observed by comparing classification performance for the attended versus unattended conditions; the stimulus displays in these two conditions are identical, but attention is directed either to the relevant category or away (Figure 2; see also Figure S10). A three-way ANOVA of ROI \times category \times attended/unattended revealed a significant main effect of category ($F(3,27) = 27.83$; $p < .0001$) and of attention ($F(1,9) = 13.8$; $p < .005$). Following up on these effects, we found a large drop in classification performance for the unattended versus attended condition. In fact, classification performance for unattended stimuli was not significantly above chance for any category in any ROI except for faces in the FFA ($t(9) = 9.93$; $p < .0001$) and

houses in the PPA ($t(9) = 2.83$; $p < .05$) and ORX (where mean performance was only 58%).

Thus, faces and houses in the FFA and PPA, respectively, were partially spared from the large drop in performance observed for other unattended categories. This dependence of attentional effects on category and ROI was supported by significant interactions of ROI by category ($F(6,54) = 19.7$; $p < .0001$), category by attended/unattended ($F(3,27) = 14.4$; $p < .0001$), and the triple interaction ($F(6,54) = 3.13$; $p < .05$). The drop in performance for unattended faces in the FFA was particularly small (from 96% to 86%; $t(9) = 2.46$; $p = .04$). This preservation of performance in the FFA was not due to an inability to direct attention away from faces—performance in ORX dropped from 79% to chance (50%) for attended versus unattended faces; this interaction of attended/unattended \times FFA/ORX on face discrimination was significant ($F(1,9) = 9.15$; $p < .05$). A corresponding preservation of above-chance performance was found for unattended houses in the PPA; the interaction of attended/unattended \times PPA/ORX was significant ($F(1,9) = 9.18$; $p < .05$).

Thus, directing attention away from an object eliminates any discriminative information in all ROIs, except for faces in the FFA and houses in the PPA and ORX. In other words, category-selective regions confer partial robustness to diverted attention.

Discussion

The main finding of this study is that information present in the spatial profile of fMRI response for common objects (e.g., cars versus shoes) is severely degraded when these objects are presented in clutter and is not above chance when attention is diverted. In contrast, information about “special” categories (faces and houses) that selectively activate particular cortical regions (the FFA and PPA, respectively) is remarkably robust to clutter and relatively preserved even when attention is diverted. Two important conclusions follow from these results. First, the distributed patterns of fMRI responses [5] are not likely to be of much use in real-world vision where images are characterized by extensive clutter. Second, category selectivity in cortex confers robustness to clutter, such that the ability to detect the presence of faces or houses based on the FFA and PPA is undiminished by the presence of another object. These findings have important implications for the utility and shortcomings of “nonpreferred” fMRI responses in neural codes in general and for the nature of cortical representations of visually presented objects in particular.

On the first point, our data empirically highlight an insight derived from computational considerations: although distributed codes have a larger representational capacity (an exponential rather than linear function of the number of units), the increased capacity comes at the cost that information is degraded when multiple representations are superimposed in the same substrate. Put another way, the pattern of weak “nonpreferred” responses can contain stimulus information [5], but much of this information may be lost when multiple objects are simultaneously present. Given that vulnerability to clutter results from overlapping representations, clutter tolerance can be achieved by reducing the overlap. In the

extreme case, grandmother cells are not vulnerable to clutter because they are not activated by anything other than your grandmother. Although the problems with such extreme local codes are well known [2], the present results serve to empirically demonstrate the shortcomings of coding schemes at the opposite extreme (i.e., highly distributed overlapping codes). Where any given perceptual system lands on the spectrum between completely local versus completely distributed coding will depend in part on how that system handles the tradeoff between representational capacity and clutter tolerance, and this in turn is likely to depend on the number of alternative stimuli that need to be coded and the amount of clutter in the natural environment where this system functions. These considerations are likely to apply across multiple domains of representation, and across multiple scales, from neurons to voxels to cortical areas.

Second, our finding that classification performance for shoes and cars is drastically impaired by clutter and diverted attention indicates important limits on the utility of distributed patterns of fMRI response as representations of objects in the real world [5, 15]. Natural scenes are substantially more cluttered than the two-object displays used in our study, and current evidence indicates that category-selective regions are not found for most object categories [11, 16]. Thus, under natural conditions, the spatial profile of fMRI response may provide robust information for only a small number of “special” categories. Additionally, note that our experiments address only a very coarse level of object classification (face/house/shoe/car), leaving open the question of whether neural codes for finer-grained discriminations (e.g., subordinate level categorizations) are robust to clutter and diverted attention.

The clutter costs revealed for “nonspecial” categories would not argue against a role for such representations in the human perceptual system, if they were also observed behaviorally. However, during the fMRI experiment, subjects discriminated exemplars of all four categories equally well in a 1-back task. Additionally, in a separate experiment, subjects performed nearly perfectly (Figure S9) when detecting a particular category (similar to the decision made by the classifiers). Thus, the information not present for nonspecial categories in the fMRI response pattern is nonetheless available to behavior, suggesting that these patterns cannot be the representations underlying the perceptual experience of categories like shoes and cars.

Although indicating important boundary conditions on the utility of distributed fMRI response patterns as codes for object category, the present findings leave some questions unanswered. First, our results argue for the relative sparing of faces and houses from the costs of clutter and diverted attention, based on fMRI response patterns in the FFA and PPA. Is there any behavioral correlate of the clutter tolerance for special categories found here? As noted above, subjects detected each of our categories nearly perfectly under the conditions of the fMRI experiment. However, given this ceiling performance, our behavioral data is not sensitive enough to pick up subtle differences in behavioral processing of different objects. Indeed, under some experimental conditions, faces and bodies appear to have

a processing advantage compared to other categories [17–22], and this advantage might constitute a behavioral correlate of the functional distinctions observed here.

Second, our conclusion that faces and houses are preserved from clutter in the fMRI analyses rests on the assumption that whoever is reading out this neural code “knows” where in the cortex to look for this information. That is, the FFA confers robustness to clutter and diverted attention for face stimuli only if the subject reads out the neural code from the FFA. If instead face classification performance is based on the union of the FFA, PPA, and ORX, then the corresponding performance is much lower in the attended (83.6% in the union ROI versus 96.2% in the FFA) and unattended (56.1% versus 85.7%) conditions. One interesting possibility is that the brain solves the problem of “knowing where to look” by simply reading out the strongest responses within a larger pattern. Consistent with this possibility, we found that faces and houses are preserved from clutter costs even when no ROI is specified in advance, if pattern information is based on the most active voxels (Figure S8).

Third, how are nonspecial categories represented in the brain, such that subjects can detect them in clutter? fMRI drastically undersamples the information present in the neural code: the data in this study cannot resolve the temporal properties of the neural response, and they provide only coarse spatial information at a grain of tens of thousands of neurons per voxel. Category information present in the responses of individual neurons will not be detectable with fMRI unless those neurons are sufficiently clustered in cortex to produce a differential response at the voxel level. Thus, it is possible that object information robust to clutter is present in the pattern of response across individual neurons but is so far not detectable with fMRI. More extensive neurophysiological investigations of the questions addressed here are warranted [23].

Finally, in the current study, as in a recent neurophysiological study [23], we investigated the effects of clutter by using just two objects presented simultaneously. Under these rather simplified clutter conditions, we already find that “standard” object categories like shoes and cars suffer large clutter costs. In more naturalistic viewing conditions, clutter costs would presumably only be greater. Ongoing work is investigating the effects of more realistic cluttered environments.

In conclusion, the present results indicate that distributed cortical codes for object category that are detectable with fMRI are vulnerable to clutter and diverted attention, suggesting limitations on the utility of such codes in natural viewing conditions. At the same time, our findings show that category-selective responses in the FFA and PPA provide substantial sparing of face and house stimuli, respectively, from the costs of clutter and diverted attention. These findings suggest that one function of category-selective neural responses may be to preserve information about the presence of biologically important categories under natural viewing conditions. Analogously, in the olfactory system of the fly, although most odors are coded by activity across multiple glomeruli [24], some biologically important odors are coded by activity in a single glomerulus that responds

exclusively to that odor [25]. More generally, our findings invite future research investigating which if any distributed cortical codes are read out behaviorally [26], whether this readout process is modulated by context and task, and whether the answers to these questions differ substantially for representations of objects at the finer grain of populations of individual neurons [27].

Experimental Procedures

Subjects

Ten healthy subjects (4 females) participated in this study. All subjects gave signed informed consent prior to the start of the experiments and had normal or corrected-to-normal vision.

Experiment Timeline

Images from four different categories (faces, houses, shoes, and cars) were used during the main experiment. Images were divided into two sets, one used on odd runs and one on even runs (30 stimuli per category in each set). The experiment followed a blocked design; in each run, subjects were presented with two blocks of each stimulus category in the isolated condition and two blocks of each pair of stimulus categories in the clutter condition (“attended” and “unattended” conditions). Subjects were instructed to press a button whenever they saw two consecutive identical images (1-back task) from a category cued at the start of each block with a 2 s instruction (e.g., “attend faces”). In each block of 20 trials, a 1-back repeat occurred twice per category. Behavioral data were collected for 8 of the 10 subjects; the response box was not functioning properly when the other 2 subjects were tested (which does not prevent analyses of BOLD patterns). In the attended and unattended conditions, all possible pairs of object categories were presented in separate blocks.

Each block was 16 s long and each stimulus was presented for 800 ms. Within each block, the images of a given category were alternately presented to the left and right of fixation as shown in Figure 1. This prevented subjects from attending to just one side of fixation during an entire block. The images were centered at 4 degrees on either side of fixation and subtended approximately 7 degrees of visual angle. Each subject was tested on 6–7 experimental runs in the scanner. The order of blocks was counterbalanced across subjects. Stimuli were presented with the Psychophysics Toolbox. Other stimulus conditions that were included to test different hypotheses will not be reported here.

Localizer Scans

In separate localizer runs, performed in the same scan sessions, subjects were presented with blocks of faces, scenes, objects (a variety of different everyday objects, including shoes and cars), and scrambled images. Scenes were used instead of houses in these localizer scans because the PPA is known to prefer scenes to houses. Subjects’ task again was a 1-back task. Three localizer runs were performed for each subject.

fMRI Data Acquisition

fMRI data was acquired on a 3T Siemens scanner at the MGH-NMR center in Charlestown, MA, for four subjects and the MIT McGovern Institute for the remaining six subjects. A Gradient Echo pulse sequence was used with a TR = 2 s and TE = 30 ms. 20 slices were collected with a 12-channel head coil. The slice thickness was 2 mm and the in-plane voxel dimensions were 1.6 × 1.6 mm. The slices were oriented roughly perpendicular to the calcarine sulcus and covered the entire temporal lobe and part of the occipital lobe. High-resolution MPRAGE anatomical images were also acquired for each subject.

fMRI Data Analysis

Data analysis was performed with FS-FAST (<http://surfer.nmr.mgh.harvard.edu>) and fROI (<http://froi.sourceforge.net>). Before statistical analysis, images were motion corrected. For the localizer runs only, the fMRI blocks were smoothed with a 3 mm full width at half maximum Gaussian kernel. Based on the data obtained during the independent localizer runs, the FFA, PPA, and object responsive (ORX)

ROIs were defined for each subject. The FFA was defined as the set of contiguous voxels in the fusiform gyrus that showed significantly stronger activation ($p < 10^{-4}$, uncorrected) to faces than to other objects. The PPA was defined as the set of voxels in the parahippocampal region that showed significantly higher activation to scenes than to other objects. The ORX ROI was the set of voxels in the ventral visual pathway that were more strongly activated to faces, objects, or scenes compared to scrambled images ($p < 10^{-4}$) with the exclusion of face-selective (all voxels with activation higher for faces than objects [$p < 10^{-3}$]) and scene-selective (all voxels with activation higher for scenes than objects [$p < 10^{-3}$]) voxels. Note that ORX is larger than LOC, which is usually defined by the contrast objects versus scrambled objects, in the lateral occipital cortex.

For the localizer scans, 3D statistical maps were calculated for each ROI by correlating the signal time course with a gamma function ($\delta = 2.25$, $\tau = 1.25$) for each voxel based on the haemodynamic response properties.

Correlation Analysis

Category information was read out from each ROI by a correlation-based technique similar to Haxby et al. [5]. The fMRI data was split in two halves based on even and odd experimental runs. The correlations between the average activation pattern of each half was computed for each pair of conditions. Based on these correlation values, object identity was predicted. For each half, the isolated blocks from the other half were used as the reference condition. To decode category information in the isolated blocks (e.g., for an isolated face block), the within- and between- category correlations were compared (i.e., face-face isolated block correlations were compared to isolated face-isolated house, isolated face-isolated shoe, and isolated face-isolated car correlations). Decoding accuracy was the proportion of within-category correlations that were larger than the between-category correlations.

To decode category information in the clutter conditions (attended/unattended), each block was compared to the isolated reference conditions. Thus, for example, to decode the face category, all the blocks in which faces appeared with another category (i.e., face-house, face-shoe, and face-car blocks) were correlated with the isolated face blocks to yield within-category correlations, and with isolated house, shoe, and car blocks to yield between-category correlations. Again, decoding accuracy was the percentage of within-category correlations larger than the between-category correlations.

To make the decoding analysis across ROIs comparable, we limited the number of voxels that each ROI could contribute to the analysis. In the analysis presented here, each ROI contributed up to 250 voxels for the analysis. Similar results were obtained when each ROI contributed up to 100, 150, or 200 voxels. To compute the analysis, each ROI was sampled 100 times and 250 voxels were randomly picked each time to participate in the correlation-based analysis. The data reported here are the average decoding performance across the 100 iterations.

The repeated-measures ANOVA analyses reported here were followed by post-hoc tests where appropriate, which were corrected for multiple comparisons by Scheffe’s method.

Eye-Tracking Control

Eye tracking was performed in a separate session with an IR iScan Camera, sampling the eye position at 240 Hz. The eye positions were calibrated at the start of each run, and approximately every 3.5 min thereafter. 5 of the 10 subjects who performed the fMRI main experiment were seated 75 cm in front of a screen and performed the experiment with the same design and stimulus sets as they had experienced in the scanner. Viewing conditions in the scanner were replicated as closely as possible during this control experiment. Subjects were not explicitly instructed to fixate but were told to perform the experiment with the same strategy they had used during the fMRI scanning session (when they had been strictly instructed to fixate).

Supplemental Data

Ten figures and Experimental Procedures are available at <http://www.current-biology.com/cgi/content/full/17/23/2067/DC1/>.

Acknowledgments

This work was supported by grant EY 13455 to N.K. We would like to thank C. Baker, P. Cavanagh, D. Dilks, P. Foldiak, J. McClelland, H. Op de Beeck, A. Pouget, R. Schwarzlose, R. VanRullen, and M. Williams for comments on the manuscript.

Received: August 25, 2007

Revised: October 18, 2007

Accepted: October 19, 2007

Published online: November 8, 2007

References

1. Willshaw, D.J. (1981). Holography, associative memory, and inductive generalization. In *Parallel Models of Associative Memory*, G.E. Hinton and J.A. Anderson, eds. (Hillsdale, NJ: Erlbaum), pp. 83–104.
2. Rumelhart, D.E., and McClelland, J.L. (1986). *Parallel distributed processing: explorations in the microstructure of cognition, Volume 1* (Cambridge: MIT Press).
3. Foldiak, P. (2002). Sparse coding in the primate cortex. In *The Handbook of Brain Theory and Neural Networks, Second Edition*, M.A. Arbib, ed. (Cambridge, MA: MIT Press), pp. 1064–1068.
4. Mel, B.W., and Fiser, J. (2000). Minimizing binding errors using learned conjunctive features. *Neural Comput.* *12*, 731–762.
5. Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., and Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* *293*, 2425–2430.
6. Spiridon, M., and Kanwisher, N. (2002). How distributed is visual category information in human occipito-temporal cortex? An fMRI study. *Neuron* *35*, 1157–1165.
7. Carlson, T.A., Schrater, P., and He, S. (2003). Patterns of activity in the categorical representations of objects. *J. Cogn. Neurosci.* *15*, 704–717.
8. O'Toole, A.J., Jiang, F., Abdi, H., and Haxby, J.V. (2005). Partially distributed representations of objects and faces in ventral temporal cortex. *J. Cogn. Neurosci.* *17*, 580–590.
9. Kanwisher, N., McDermott, J., and Chun, M.M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J. Neurosci.* *17*, 4302–4311.
10. Epstein, R., and Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature* *392*, 598–601.
11. Baker, C.I., Hutchison, T.L., and Kanwisher, N. (2007). Does the fusiform face area contain subregions highly selective for nonfaces? *Nat. Neurosci.* *10*, 3–4.
12. Haynes, J.D., and Rees, G. (2006). Decoding mental states from brain activity in humans. *Nat. Rev. Neurosci.* *7*, 523–534.
13. Norman, K.A., Polyn, S.M., Detre, G.J., and Haxby, J.V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn. Sci.* *10*, 424–430.
14. Vapnik, V.N. (1998). *Statistical Learning Theory* (New York: Wiley).
15. Cohen, J.D., and Tong, F. (2001). Neuroscience. The face of controversy. *Science* *293*, 2405–2407.
16. Downing, P.E., Chan, A.W., Peelen, M.V., Dodds, C.M., and Kanwisher, N. (2006). Domain specificity in visual cortex. *Cereb. Cortex* *16*, 1453–1461.
17. Downing, P.E., Bray, D., Rogers, J., and Childs, C. (2004). Bodies capture attention when nothing is expected. *Cognition* *93*, B27–B38.
18. Kirchner, H., and Thorpe, S.J. (2005). Ultra-rapid object detection with saccadic eye movements: visual processing speed revisited. *Vision Res.* *46*, 1762–1776.
19. Thorpe, S., Crouzet, S., Kirchner, H., and Fabre-Thorpe, M. (2006). Ultra rapid face detection in natural images: implications for computation in the visual system. In *First French Conference on Computational Neurosciences (Abbaye des Premontres, PontaMousson, France)*, pp. 124–127.
20. Awh, E., Serences, J., Laurey, P., Dhaliwal, H., van der Jagt, T., and Dassonville, P. (2004). Evidence against a central bottleneck during the attentional blink: multiple channels for configural and featural processing. *Cognit. Psychol.* *48*, 95–126.
21. Einhauser, W., Koch, C., and Makeig, S. (2007). The duration of the attentional blink in natural scenes depends on stimulus category. *Vision Res.* *47*, 597–607.
22. Ro, T., Russell, C., and Lavie, N. (2001). Changing faces: a detection advantage in the flicker paradigm. *Psychol. Sci.* *12*, 94–99.
23. Zoccolan, D., Cox, D.D., and DiCarlo, J.J. (2005). Multiple object response normalization in monkey inferotemporal cortex. *J. Neurosci.* *25*, 8150–8164.
24. Wang, J.W., Wong, A.M., Flores, J., Vosshall, L.B., and Axel, R. (2003). Two-photon calcium imaging reveals an odor-evoked map of activity in the fly brain. *Cell* *112*, 271–282.
25. Suh, G.S., Wong, A.M., Hergarden, A.C., Wang, J.W., Simon, A.F., Benzer, S., Axel, R., and Anderson, D.J. (2004). A single population of olfactory sensory neurons mediates an innate avoidance behaviour in *Drosophila*. *Nature* *431*, 854–859.
26. Williams, M.A., Dang, S., and Kanwisher, N.G. (2007). Only some spatial patterns of fMRI response are read out in task performance. *Nat. Neurosci.* *10*, 685–686.
27. Reddy, L., and Kanwisher, N. (2006). Coding of visual objects in the ventral stream. *Curr. Opin. Neurobiol.* *16*, 408–414.