

The Neural Basis of the Behavioral Face-Inversion Effect

Galit Yovel,^{1,2,*} and Nancy Kanwisher¹

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

Summary

Two of the most robust markers for “special” face processing are the behavioral face-inversion effect (FIE)—the disproportionate drop in recognition of upside-down (inverted) stimuli relative to upright faces—and the face-selective fMRI response in the fusiform face area (FFA). However, the relationship between these two face-selective markers is unknown. Here we report that the behavioral FIE is closely associated with the fMRI response in the FFA, but not in other face-selective or object-selective regions. The FFA and the face-selective region in the superior temporal sulcus (f_STS), but not the occipital face-selective region (OFA), showed a higher response to upright than inverted faces. However, only in the FFA was this fMRI-FIE positively correlated across subjects with the behavioral FIE. Second, the FFA, but not the f_STS, showed greater neural sensitivity to differences between faces when they were upright than inverted, suggesting a possible neural mechanism for the behavioral FIE. Although a similar trend was found in the occipital face area (OFA), it was less robust than the FFA. Taken together, our data suggest that among the face-selective and object-selective regions, the FFA is a primary neural source of the behavioral FIE.

Results and Discussion

The present study tested the role of each of the three face-selective regions (FFA, OFA, and f_STS), plus an object-selective region called the lateral occipital complex or LOC [1, 2], in the behavioral face-inversion effect. We predicted that cortical region(s) underlying the behavioral FIE should show a close association between fMRI responses and behavioral discrimination of upright and inverted faces. In particular, our first prediction was that the difference in fMRI response to upright and inverted faces (fMR-FIE) will be correlated across subjects with the difference in performance level between upright and inverted faces (behavioral FIE). Second, a cortical region involved in the behavioral FIE would be expected to show greater sensitivity to differences between faces when they are upright than inverted. That is, a region that is involved in the FIE will show a greater fMR-adaptation effect [3, 4] when faces are upright than inverted.

Behavioral and fMRI measures of the FIE were collected in two experiments. In one experiment ($n = 21$), we used an event-related design in which upright- and inverted-face trials were presented randomly within the same scan. In a second experiment ($n = 14$), we used a blocked design where upright- and inverted-face trials were presented in separate blocks.

Behavioral Results

In both experiments, subjects performed a task during the main fMRI scans in which they were asked whether the two sequentially presented faces in each trial were the same or different; the two faces in each trial were either both upright or both inverted (see Figure 1A). Behavioral data showed a significant drop in accuracy for inverted compared to upright faces (i.e., a behavioral FIE): event-related experiment: upright: 82%, inverted: 75%; $t(15) = 5.37$, $p < .0001$; blocked-design experiment: upright: 77%, inverted: 69%, $t(13) = 4.29$, $p < .001$.

Functional MRI Analyses

Localizer Results

Each subject participated in a blocked localizer scan (see Figure 1B) in the same scanning session as the main experiment, in which they viewed faces, objects, and scrambled objects. A comparison of faces > objects ($p < 10^{-4}$, uncorrected) served to functionally identify in each subject the following regions of interest (ROIs) across the two experiments: FFA (in 100% of subjects), OFA (in 89% of subjects), and f_STS (in 74% of subjects). The LOC was identified in 100% of subjects with a comparison of objects > scrambled objects ($p < 10^{-4}$, uncorrected).

In a few of the subjects in the event-related experiment, the face-selective areas that were revealed by the localizer did not show a higher response to faces than nonfaces (i.e., chairs) in the experimental task (see Stimuli section), which suggests that the localizer-defined face-selective regions were not reliable. Thus, only subjects that showed higher responses to faces in the localizer-defined face-selective regions were included in the analyses ($n = 17$ for FFA, $n = 14$ for OFA, and $n = 11$ for f_STS). All the subjects in the blocked-design task showed higher response to faces than nonfaces (houses) in the face-selective regions that were defined by the localizer.

Main Experimental Task

Having identified the ROIs in each subject based on the localizer data, we then calculated the magnitude of the BOLD response in each of these ROIs in each of the conditions of the main experiments. These data were used to test whether each of the four ROIs is involved in generating the behavioral FIE. Analyses showed no difference in the pattern of response between the right and left hemispheres for each of the ROIs. Therefore, all analyses are based on a pooled analysis in which right and left hemisphere voxels are combined in each ROI.

*Correspondence: galit@freud.tau.ac.il

²Present address: Department of Psychology, Tel Aviv University, Tel Aviv 69978, Israel.




Figure 1. Procedure of the Main and Localizer Tasks

(A) Experimental task. In both experiments, subjects were presented with a sequential matching face task. The upright and inverted trials were mixed in the event-related study and in separate blocks in the blocked-design study. The stimuli presented in the figure are sample stimuli from the event-related experiment.

(B) LocalizerTask. Subjects were presented with 16 s blocks of faces, objects, and scrambled images of objects (bodies and scenes were also presented but are not included in the analysis). Each block included 20 stimuli, and subjects were instructed to press a key whenever two images were repeated consecutively (one-back task). We localized face-selective regions (face > object) and object-selective regions (object > scrambled objects). For more details about the localizer, see [21].

1. Functional MRI Face-Inversion Effect. Assessment of the inversion effect in each ROI showed a significantly higher response to upright than inverted faces in the FFA [$t(30) = 4.67$, $p < .0001$] and f_STS [$t(21) = 3.89$, $p < .001$], but not in the OFA [$t(27) = .98$, $p > .33$], and an opposite effect of a higher response to inverted than upright faces in LOC [$t(34) = 2.31$, $p < .03$]. ANOVA with area and orientation as repeated measures showed that the inversion effect was significantly larger in the FFA than the OFA [$F(1,26) = 6.25$, $p < .02$] and significantly different from the LOC [$F(1,28) = 32.31$, $p < .00001$] but not from the f_STS (see [Supplemental Data](#) available with this article online for an omnibus ANOVA).

Thus, the FFA and f_STS showed a higher overall response to upright than inverted faces, whereas the OFA showed no significant difference and the LOC showed an opposite effect of a higher response to inverted than upright faces. These analyses might seem to implicate the f_STS and FFA (and perhaps also the LOC) in the behavioral face-inversion effect. However, the mere existence of an fMRI-FIE need not imply that it is necessarily related to the behavioral FIE. We therefore went on to conduct a stronger test of the relationship between the fMRI-FIE and the behavioral FIE.

2. A Correlation across Subjects between fMRI and Behavioral Measures of the Face-Inversion Effect. We




Figure 2. The Response of Each of the Three Face-Selective Regions and the Object-Selective Region to Upright and Inverted Faces
Of the three face-selective regions (FFA, OFA, and f_STS), the FFA and f_STS showed a higher response to upright than inverted faces (fMRI-FIE). The OFA showed no difference in response to upright and inverted faces and the object-selective region, LOC, showed an inverted fMRI-FIE, with

tested whether the fMR-FIE was correlated across subjects with the behavioral FIE as follows. We calculated a fMRI-FIE score (BOLD response for upright – inverted) for each ROI in each subject. In addition, a behavioral FIE score was also calculated for each subject (accuracy for upright – inverted): in the face-matching tasks. We then measured the correlation across subjects between these behavioral and neural FIE scores (see Figure 2, right column), separately for each ROI. To increase the power of the correlation analysis, we included data from both the event-related and the blocked-design experiments in this analysis.

The correlation analyses showed that only the FFA-FIE was positively correlated with the behavioral FIE [$r(28) = .50, p < .005$]. Interestingly, the f_STS, which showed a higher response to upright than inverted faces in mean response, did not show a correlation between the f_STS-FIE and the behavioral FIE [$r(18) = .06, p > .79$]. This finding is consistent with the hypothesis that the f_STS is specialized for representing dynamic aspects of facial information such as expression and gaze [5, 6], not identity information [7], which was measured with our behavioral task. The OFA, which showed no difference between the response to upright versus inverted faces in group means, also failed to show a correlation between the OFA-FIE and the behavioral FIE [$r(26) = -.004, p > .98$]. These findings are consistent with a recent report, which suggests that the OFA is sensitive to physical aspects of the face stimulus rather than to face identity [8].

Finally, previous studies suggested that the higher response to inverted than upright faces in house-selective areas (i.e., areas that showed higher response to houses than faces), which may partly overlap with the ventral part of LOC, might be the neural source of the behavioral FIE [9, 10]. In contrast to this hypothesis, we found that the higher response to inverted than upright faces in the object-selective LOC (LOC-FIE) was not correlated with the behavioral FIE [$r(32) = .05, p > .77$]. Notably, the lack of correlation between the behavioral FIE and the OFA-FIE or f_STS-FIE can not be due to low reliability of the fMRI-FIE measures from these two face-selective regions: the f_STS-FIE and OFA-FIE were strongly correlated with each other [$r(19) = .70, p < .01$], which suggest that the FIE in these regions were highly reliable and that their lack of correlation with the behavioral FIE reflects lacks of association rather than low reliability. The zero order correlations between the fMR-FIE of all face-selective regions are reported in Table 1 (see Supplemental Data for correlational analyses with a normalized FIE measure and a multiple regression analysis that assesses that contribution of each face-selective region to the behavioral-FIE).

In summary, our analyses indicate that among the face-selective and object-selective regions, only the FFA-FIE is related to the behavioral FIE. Next, we used event-related fMRI adaptation to test whether the fMR

Table 1. The Zero-Order Correlations between the fMRI Face-Inversion Effects of the Three Face-Selective ROIs and the Object-Selective ROI

	FFA-FIE	OFA-FIE	STS-FIE	LOC-FIE
FFA-FIE		0.40 n = 26 p < .05	0.52 n = 21 p < .02	0.16 n = 30 p > .38
OFA-FIE			0.70 n = 19 p < .001	0.23 n = 27 p > .25
STS-FIE				0.17 n = 22 p > .45

Correlation was calculated by the following formula: % signal change_upright – % signal change_inverted.

response is more sensitive to differences between faces when they are upright than when they are inverted.

3. Neural Discrimination of Identity Information in Upright and Inverted Faces. The most straightforward account of the behavioral FIE is that neural populations in face-processing regions discriminate individual identity of faces better when they are upright than when they are inverted. To assess whether each ROI is more sensitive to differences between faces when they are presented upright than inverted, we measured the fMR adaptation to face identity for upright and inverted faces by separately analyzing the fMRI response on trials in which two different faces were presented from trials in which the same face was presented twice. A higher response on different than same trials in a given region would indicate discrimination of the two different stimuli by neural populations within that region [4]. An interaction of orientation and adaptation (different/same), which reflects a higher response on different than same trials for upright faces, but not for inverted faces, was found for the FFA [$F(1,16) = 7.93, p = .01$] and OFA [$F(1,13) = 11.04, p < .01$], but not the f_STS or LOC (see Figure 3). These findings are indicative of better neural discrimination of upright than inverted faces in the FFA and OFA. (See Supplemental Results for analyses of simple effects within regions, interaction across regions, and analyses based on only correct responses.)

In sum, although the OFA-FIE showed similar response to upright and inverted faces and was not correlated with the behavioral FIE, the stronger adaptation to upright than inverted faces in the OFA suggests that it does treat upright and inverted faces differently. Thus, our findings with respect to the role of the OFA in the behavioral FIE are not conclusive. Importantly, both the correlation and the adaptation findings clearly suggest that the f_STS do not play a role in the behavioral FIE in an identity discrimination task, which is consistent with the suggestion that the f_STS is involved in dynamic (e.g., gaze, expression), rather than static (e.g., identity) aspects of face processing [11]. Finally, in contrast to

a higher response to inverted than upright faces (middle column). Error bars represent the standard error of the difference between upright and inverted faces. Note that the scale of the y axis is not the same for all three regions because the response in the OFA (~1.1) to faces is much larger than the response of the f_STS and LOC (~.4). For each ROI, we calculated the correlation across subjects between the fMRI-FIE and the behavioral FIE. The correlation is represented in a scatterplot for each ROI (right column) in which each subject is represented by a point, with that subject's fMRI-FIE score for each ROI on the x axis and the same subject's behavioral FIE score on the y axis. Only the FFA-FIE was correlated with the behavioral FIE.

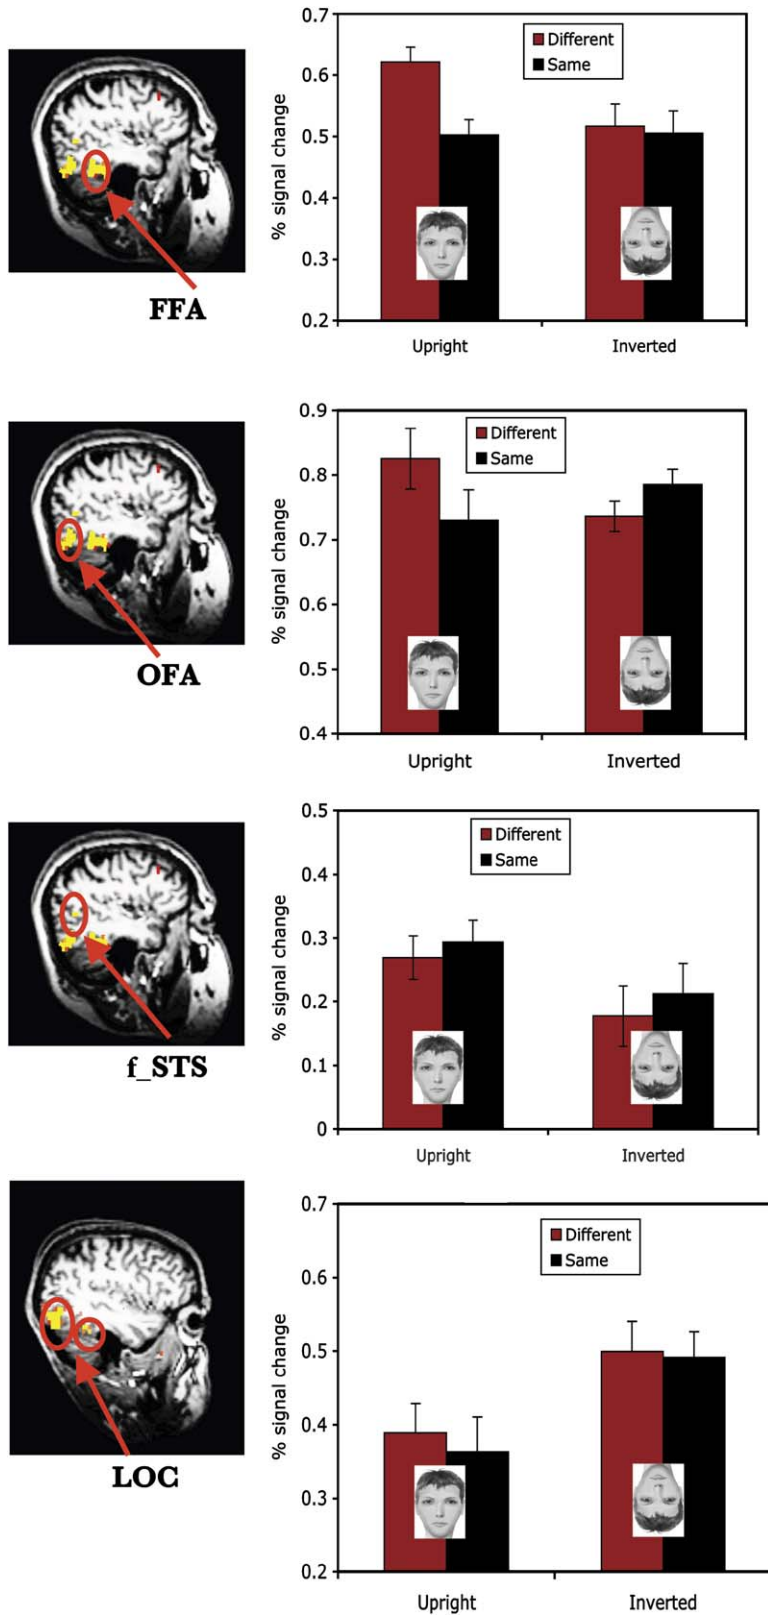


Figure 3. fMRI Adaptation to Upright and Inverted Faces in Each of the Face-Selective Regions and the Object-Selective Region

The fMRI response of each of the face-selective (FFA, OFA, and f_STS) and object-selective (LOC) regions to pairs of different versus same faces that were presented upright or inverted. A higher response to different than same pairs indicates neural discrimination between faces. The FFA and OFA showed better discrimination between upright than between inverted faces. The f_STS and LOC did not discriminate either upright or inverted faces. Error bars represent the standard error of the difference between the response to different and same trials. Note that the scale of the y axis is not the same for all three regions because the response in the OFA (~.8) to faces is much larger than the response of the f_STS and LOC (~.3).

prior claims [9, 10] that the higher response to inverted than upright faces in the LOC might be related to the behavioral FIE [9, 10], both the correlation and the adaptation data in our study suggest that the LOC plays no role in the behavioral FIE.

Conclusion

Our findings indicate that among the face-selective and object-selective regions, the FFA is a primary neural source of the behavioral FIE. First, the FFA showed a higher response to upright than inverted faces (FFA-

FIE) that was positively correlated with the behavioral FIE across subjects. Second, the FFA was more sensitive to differences between faces when they were presented upright than when presented upside-down.

These findings are important for three reasons. First, they resolve a long-standing mystery of the relationship between the two most well-established markers of face processing, the behavioral FIE and the FFA, providing strong evidence that it is neural representations in the FFA that are read out when subjects perform discriminations on upright faces [8, 12] but not inverted faces. Second, our findings provide a possible explanation of the behavioral FIE in neural terms: the lower behavioral sensitivity to differences between inverted faces compared to upright faces apparently results at least in part from sharper neural tuning for upright than inverted faces. Third, our results show clear functional dissociations between the FFA and other cortical regions engaged in face and object perception.

Four previous reports that examined the fMRI response to upright and inverted faces reported either no FIE in the FFA or a weak effect, whereas our study revealed a robust FIE in the FFA. Although the precise reasons for these varying results are not yet clear, we speculate that the fMRI-FIE is larger and more robust for tasks that generate a strong behavioral face inversion effect, which was not included in most previous reports.

What Is the Nature of the Representation of Upright Faces?

By linking the FFA to the classic behavioral hallmark of special face processing, the behavioral FIE, the present results open the door to a deeper investigation of the nature of the representations underlying face recognition. Here again behavioral work provides important clues. Numerous behavioral studies have indicated that what is “special” about the processing of upright but not inverted faces is that only upright faces are processed “holistically,” such that face parts are processed interactively rather than independently (for review, see [13]). Indeed, a recent fMRI study that used the composite face task [14], which provides a behavioral measure of holistic processing, found a neural correlate of this effect in the FFA (C. Schiltz and B. Rossion, 2005, *Vision Sci. Soc.*, abstract).

However, fMRI investigations are of course severely limited in both their spatial and temporal resolution, leaving open many questions about the neural mechanisms responsible for the differences in the FFA response to upright and inverted faces. Do the inversion effects reported here reflect the operation of a single neural population in which most neurons respond more strongly to upright than inverted faces [15, 16], or does the FFA contain distinct neural populations, one responsive only to upright faces and another responsive only to inverted faces? Evidence for at least partly separate representations for upright and inverted faces comes from a behavioral study that found opposite figural aftereffects simultaneously induced on upright and inverted faces [17]. Another important dimension of the FIE unaddressed here concerns the temporal profile of the processing of upright and inverted faces. Event-related potential studies in humans [18] and single-cell recordings in macaques [19] have found that latencies of neural responses to inverted faces are delayed

compared to upright faces. It will be important for future studies to test whether the fMRI signature of the FIE reported here is associated with the differences in response to upright and inverted faces that are found in electrophysiological studies.

To sum up, the present findings have advanced our understanding of the neural basis of face processing by clearly linking fMRI and behavioral markers of specialized face processing (namely, the FFA and the behavioral FIE), by functionally dissociating the FFA from the OFA, F_STS, and LOC, and by suggesting a possible neural mechanism underlying the behavioral FIE (namely, sharper neural tuning for upright than inverted faces). Continued progress in understanding how faces are represented in the brain is likely to result from further fMRI investigations motivated by the rich behavioral literature on face processing.

Experimental Procedures

Event-Related FIE Experiment

Stimuli

Eighteen different faces were generated by FACE 3.0 software. There were six sets of three faces that shared the same hair but differed in internal facial information and face outline. In each trial, the pairs of face stimuli shared the same hair, to assure that subjects use facial information, rather than hair per se, to identify the stimuli. The face stimuli subtended 4.5° (width) × 7° (length) of visual angle. Chair stimuli were also included in the experimental task, which allowed us to determine for each subject whether the experimental task showed selectivity to faces (faces > chairs) in the voxels that were identified as face selective by the localizer (see [Localizer Results](#)).

Procedure

The experiment consisted of five runs of the experimental task and five runs of the localizer, which were presented interleaved with the experimental task.

Experimental Task. Pairs of face or chair stimuli were presented sequentially either upright or inverted in a randomized order that was optimized for the extraction of the hemodynamic response in an event-related fast presentation design. Each trial lasted 2 s. The first and second stimuli were presented for 250 ms, with an interstimulus interval of 500 ms. The next trial was presented 1 s after the disappearance of the second stimulus. Five runs of the experimental task were included. Each run included 24 trials of each condition. In half of the trials, the stimuli were identical, and in the other half, they were different. Each stimulus was presented equally often in the identical and different conditions, so any difference in the fMRI response between conditions must be due to the relationship between the two stimuli in a pair, not the individual stimuli themselves. Same-different responses were made with a key press. Twenty-four blank trials were presented intermixed with the experimental trials. Each scan lasted 240 s. Subjects made a same/different response on each trial. Behavioral responses were collected to obtain a behavioral measure of the face-inversion effect. For information about the localizer scan, see [Figure 1B](#) and [21].

fMRI Data Acquisition

Scanning was done on a 3T Allegra scanner at the MGH/MIT/HMS Athinoula A. Martinos Center for Biomedical research in Charlestown, MA. A head coil and a Gradient Echo pulse sequence with TR 2 s, TE 30 ms; flip angle 90° were used. Twenty-eight 4 mm thick slices that covered the entire brain were oriented parallel to the temporal lobe.

fMRI Data Analysis

Data were analyzed separately for each subject with FS-fast (<http://surfer.nmr.mgh.harvard.edu/>). Motion correction was conducted prior to data analysis with the AFNI motion correction algorithm [20] to align all the images to the first time image of the first run.

Blocked Design (Localizer). Data from the localizer were spatially smoothed with a Gaussian filter (full width half maximum = 5 mm) and were used to define the ROIs separately for each subject. A γ function with $\delta = 2.25$ and $\tau = 1.25$ was used to estimate the

hemodynamic response (HDR) for each condition in the localizer scans.

Event-Related Design (Experimental Task). A deconvolution analysis was used for the analysis of the event-related experimental task to extract the HDR for each same and different pair of stimuli in the event-related adaptation task, and no assumption was made on the shape of the HDR. The data from the experimental runs were not smoothed. Statistical analysis was performed on the peak of the hemodynamic response. The time points used as the peak were determined in a hypothesis-neutral fashion based on the shape of the response in each area in each subject to the average of all the stimulus conditions; the peak time points included either the 4th TR (6–8 s) or the average of the 4th and 5th TRs, when the peak lasted more than one TR.

Blocked FIE Experiment

Stimuli

The experimental task included upright and inverted face stimuli and upright house stimuli. Here we focus only on the face stimuli. The face stimuli subtended 2.5° (width) × 4° (length) of visual angle. The face stimuli differed by either spacing among the parts or the shape of parts and were presented in different blocks in upright or inverted orientation (the mean response of the FFA across subjects to the part and spacing condition for upright and inverted faces are reported in [21]). Our findings showed no difference in fMRI response and in fMRI-FIE to the spacing and part face stimuli. Therefore, for the purpose of the investigation of the FIE, we averaged the response across the two face sets.

Procedure

The experiment included five runs of a localizer (see Figure 1B) and six runs of the experimental task. Upright and inverted face stimuli were presented in separate blocks. Each block included 10 trials in which a pair of faces that were either identical (on 50% of trials) or different were presented sequentially. Subjects performed a same-different discrimination task on each pair of stimuli by pressing one key for a “same” response and another key for a “different” response. The order of the upright and inverted blocks was counter-balanced across scans.

fMRI Data Acquisition and Analysis

See methods for event-related experiment.

Supplemental Data

Supplemental Data include one table and Supplemental Results and can be found with this article online at <http://www.current-biology.com/cgi/content/full/15/24/2256/DC1/>.

Acknowledgements

We would like to thank Chris Baker, Brad Duchaine, Johannes Haushofer, Margaret Livingston, Elinor McKone, Hans Op de Beeck, Rachel Robbins, Rebecca Schwarzlose, and Doris Tsao for their comments on the manuscript. We also thank Ming-fai Fong for help with data analyses. This research was supported by NIH grants 66696 and EY13455 to N.K., by the National Center for Research Resources (P41-RR14075, R01 RR16594-01A1, and the NCRR BIRN Morphometric Project BIRN002), and by the Mental Illness and Neuroscience Discovery (MIND) Institute.

Received: August 6, 2005

Revised: October 28, 2005

Accepted: October 31, 2005

Published: December 19, 2005

References

1. Malach, R., Reppas, J.B., Benson, R.R., Kwong, K.K., Jiang, H., Kennedy, W.A., Ledden, P.J., Brady, T.J., Rosen, B.R., and Tootell, R.B. (1995). Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proc. Natl. Acad. Sci. USA* 92, 8135–8139.
2. Kanwisher, N., Chun, M.M., McDermott, J., and Ledden, P.J. (1996). Functional imaging of human visual recognition. *Cogn. Brain Res.* 5, 55–67.
3. Grill-Spector, K., and Malach, R. (2001). fMR-adaptation: a tool for studying the functional properties of human cortical neurons. *Acta Psychol. (Amst.)* 107, 293–321.
4. Kourtzi, Z., and Kanwisher, N. (2001). Representation of perceived object shape by the human lateral occipital complex. *Science* 293, 1506–1509.
5. Puce, A., Allison, T., Bentin, S., Gore, J.C., and McCarthy, G. (1998). Temporal cortex activation in humans viewing eye and mouth movements. *J. Neurosci.* 18, 2188–2199.
6. Hooker, C.I., Paller, K.A., Gitelman, D.R., Parrish, T.B., Mesulam, M.M., and Reber, P.J. (2003). Brain networks for analyzing eye gaze. *Brain Res. Cogn. Brain Res.* 17, 406–418.
7. Hoffman, E.A., and Haxby, J.V. (2000). Distinct representations of eye gaze and identity in the distributed human neural system for face perception. *Nat. Neurosci.* 3, 80–84.
8. Rotshtein, P., Henson, R.N., Treves, A., Driver, J., and Dolan, R.J. (2005). Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. *Nat. Neurosci.* 8, 107–113.
9. Haxby, J.V., Ungerleider, L.G., Clark, V.P., Schouten, J.L., Hoffman, E.A., and Martin, A. (1999). The effect of face inversion on activity in human neural systems for face and object perception. *Neuron* 22, 189–199.
10. Rossion, B., and Gauthier, I.I. (2002). How does the brain process upright and inverted faces? *Behav. Cogn. Sci. Rev.* 1, 63–75.
11. Haxby, J.V., Hoffman, E.A., and Gobbini, M.I. (2000). The distributed human neural system for face perception. *Trends Cogn. Sci.* 4, 223–233.
12. Grill-Spector, K., Knouf, N., and Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nat. Neurosci.* 7, 555–562.
13. Tanaka, J.W., and Farah, M. (2003). The holistic representation of faces. In *Perception of Faces, Objects and Scenes: Analytic and Holistic Processes*, M.A. Peterson and G. Rhodes, eds. (New York: Oxford University Press), pp. 53–74.
14. Young, A.W., Hellawell, D., and Hay, D.C. (1987). Configurational information in face perception. *Perception* 16, 747–759.
15. Perrett, D.I., Oram, M.W., and Ashbridge, E. (1998). Evidence accumulation in cell populations responsive to faces: an account of generalisation of recognition without mental transformations. *Cognition* 67, 111–145.
16. Tanaka, K., Saito, H., Fukada, Y., and Moriya, M. (1991). Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *J. Neurophysiol.* 66, 170–189.
17. Rhodes, G., Jeffery, L., Watson, T.L., Jaquet, E., Winkler, C., and Clifford, C.W. (2004). Orientation-contingent face aftereffects and implications for face-coding mechanisms. *Curr. Biol.* 14, 2119–2123.
18. Bentin, S., Allison, T., Puce, A., Perez, E., and McCarthy, G. (1996). Electrophysiological studies of face perception in humans. *J. Cogn. Neurosci.* 8, 551–565.
19. Perrett, D.I., Smith, P.A., Potter, D.D., Mistlin, A.J., Head, A.S., Milner, A.D., and Jeeves, M.A. (1985). Visual cells in the temporal cortex sensitive to face view and gaze direction. *Proc. R. Soc. Lond. B. Biol. Sci.* 223, 293–317.
20. Cox, R.W., and Jesmanowicz, A. (1999). Real-time 3D image registration for functional MRI. *Magn. Reson. Med.* 42, 1014–1018.
21. Yovel, G., and Kanwisher, N. (2004). Face perception: domain specific not process specific. *Neuron* 44, 889–898.