RESPONSE PROPERTIES OF THE HUMAN FUSIFORM FACE AREA

Frank Tong

Harvard University, Cambridge and Massachusetts General Hospital Nuclear Magnetic Resonance Center, Charlestown, USA

Ken Nakayama

Harvard University, Cambridge, USA

Morris Moscovitch

University of Toronto, Canada

Oren Weinrib

Harvard University, Cambridge, USA

Nancy Kanwisher

Massachusetts Institute of Technology, Cambridge and Massachusetts General Hospital Nuclear Magnetic Resource Centre, Charlestown, USA

We used functional magnetic resonance imaging to study the response properties of the human fusiform face area (FFA: Kanwisher, McDermott, & Chun, 1997) to a variety of face-like stimuli in order to clarify the functional role of this region. FFA responses were found to be (1) equally strong for cat, cartoon and human faces despite very different image properties, (2) equally strong for entire human faces and faces with eyes occluded but weaker for eyes shown alone, (3) equal for front and profile views of human heads, but declining in strength as faces rotated away from view, and (4) weakest for nonface objects and houses. These results indicate that generalisation of the FFA response across very different face types cannot be explained in terms of a specific response to a salient facial feature such as the eyes or a more general response to heads. Instead, the FFA appears to be optimally tuned to the broad category of faces.

INTRODUCTION

Numerous sources of evidence suggest that primate brains have special-purpose neural machinery that is selectively involved in the perception of faces. Physiological measurements, especially single-unit recordings in macaques and event-related potentials in humans, provide some of the richest sources of evidence on the specificity of these systems. However, these techniques do not allow us to quantify responses from specific regions of the human brain. The goal of the present effort was to provide a detailed characterisation of the response properties of a region of human extrastriate cortex called the fusiform face area (Kanwisher, McDermott, & Chun, 1997). We begin with a brief outline of the neurophysiological evidence for face-specific neural systems.

Early evidence that there may be specialised neural regions for face perception came from cases of prosopagnosia, the selective loss of face recognition abilities in patients with focal brain damage (e.g. Bodamer, 1947; Meadows, 1974). However, the first neurophysiological evidence of such specialisation came from the discovery of face-selective cells in the temporal cortex of macaques (Gross, Roche-Miranda, & Bender, 1972). Since this pioneering work, many studies have demonstrated that "face cells" may be tuned to certain facial attributes such as the identity (Yamane, Kaji, & Kawano, 1988), expression (Hasselmo, Rolls, & Baylis, 1989), viewpoint (Perrett et al., 1991), or parts of a face (Perrett, Rolls, & Caan, 1982; Yamane et al., 1988).

Intracranial recordings from the human temporal lobes and hippocampus (carried out for presurgical planning purposes) have also revealed individual neurons that respond selectively to faces, particular facial expressions, or gender (Fried, MacDonald, & Wilson, 1997; Heit, Smith, & Halgren, 1988; Ojemann, Ojemann, & Lettich, 1992). Evoked potentials recorded from strip electrodes implanted on the surface of the human brain have revealed distinct regions in the fusiform and inferotemporal gyri that produce face-specific N200 responses to faces but not to cars, butterflies, or other control stimuli (Allison et al., 1994). Furthermore, electrical stimulation of these regions frequently produced a temporary inability to name famous faces, suggesting that these cortical regions are not only engaged by, but necessary for, face recognition.

Although intracranial recordings provide impressive evidence for anatomically restricted responses to faces, with this technique it is difficult to collect enough data to clearly establish the face-specificity of the responses. Further, one can never be sure that the organisation of epileptic or damaged brains resembles that of the normal population. Noninvasive scalp recordings from normal subjects using event-related potentials (ERP) or magnetoencephalography (MEG) can circumvent these problems, and indeed face-selective responses have been reported in several studies (Jeffreys, 1989, 1996; Sams, Hietanen, Hari, Ilmoniemi, & Lounasmaa, 1997).

ERP and MEG studies provide excellent temporal resolution but incomplete information about the anatomical source of the signal. By contrast, recently developed neuroimaging techniques provide a method of localising functional signals with high spatial precision and therefore provide a critical new perspective. A large number of studies have demonstrated activation of human ventral extrastriate cortex during viewing of faces (e.g. Haxby et al., 1994; Puce, Allison, Asgari, Gore, & McCarthy, 1996; Sergent, Ohta, & MacDonald, 1992). Although these earlier studies did not attempt to establish the selectivity of these responses, two recent functional magnetic resonance imaging (fMRI) studies have addressed this question. McCarthy, Puce, Gore, and Allison (1997) found that a discrete region in the right fusiform gyrus responded preferentially to faces as compared to flowers or common objects. These findings agreed with other anatomical evidence showing that lesions in prosopagnosic patients (Meadows, 1974) and intracranial face-specific responses (measured preoperatively) in epileptic patients (Allison et al., 1994) frequently involve the fusiform gyrus. Kanwisher et al. (1997) demonstrated that this region, called the fusiform face area (FFA), responded in a highly selective fashion to faces as compared to objects, houses, scrambled faces, or human hands, even when subjects performed a demanding matching task that required attention to both face and nonface stimuli.

Kanwisher et al.'s results suggest that the FFA response is unlikely to be explained in terms of differences in the low-level properties of the stimuli, a general response to anything animate or human, or a tendency for subjects to attend more to faces than nonfaces during passive viewing tasks. However, they tell us little about what aspects of a face are responsible for activating the FFA. For example, is the FFA specifically tuned to human faces or does it respond more broadly, such that any type of face or the mere presence of a head will fully activate it? Does FFA activity reflect a response to the configuration of the face alone, without detailed informa-

tion about face parts, or conversely are facial features sufficient to activate the FFA without information about the configuration of the face? By studying the response of the FFA to a variety of face-like stimuli, we hoped to elucidate the functional role of this area and address whether the FFA is primarily involved in face perception or some other type of processing such as gaze perception (Heywood & Cowey, 1992) or head recognition (Sinha & Poggio, 1996).

To address these questions, the present experiments incorporated several critical design features. For each experiment, we first functionally localised each subject's FFA on independent localiser scans conducted in the same session (see General Methods, also Kanwisher et al., 1997). Examples of the localised FFA of two subjects are shown in near-axial and near-coronal slices in Plate 3 of the colour section situated between pages 160 and 161 (please click here to see Plate 3). Such individual localisation was crucial because the FFA can vary considerably in its anatomical location and spatial extent across subjects (Kanwisher et al., 1997). By specifying our region of interest in advance, we could objectively measure the magnitude of FFA responses on experimental scans without having to correct for multiple comparisons across all scanned voxels. In each experiment, some subjects viewed stimuli passively while others performed a "one-back" matching task, which obligated them to attend to all stimuli regardless of inherent interest. To the extent that similar data were obtained in the two tasks, we could infer that the results did not reflect confounding differences in the engagement of visual attention by different stimuli.

The most important methodological advance in this study involved extending beyond simple pairwise comparisons to provide a richer description of the response profile of the FFA across a range of visual stimuli. It is well known that visual neurons tuned to simple or complex stimuli do not show all-or-none responses, but instead show a graded distribution of responses that peak around the optimal stimulus. By including four different stimulus conditions within each experimental scan, we were able to quantify the magnitude of FFA

response to each stimulus class relative to others in the same scan. Each scan contained two reference conditions: (1) an optimal stimulus condition consisting of front-view human faces that are known to produce strong FFA responses, and (2) a nonoptimal stimulus condition consisting of either nonface objects or houses that produce weak FFA responses, typically less than half the magnitude found for faces. Each scan also contained two new stimulus conditions of interest which could then be compared to these optimal and nonoptimal reference conditions. Thus we could determine whether each new stimulus category produced a weak FFA response no greater than those found for objects or houses, an intermediate level response, or an optimal FFA response as strong as those found for front-view human faces.

GENERAL METHODS

Subjects

Eighteen healthy normal adults (eight women), ages 18–39, volunteered or participated for payment in one to three of the following fMRI experiments. All subjects had normal or corrected-to-normal vision and gave informed written consent to participate in the study. Data from four subjects were discarded because of artefacts caused by excess head motion (two men) or because the FFA was not successfully localised on independent scans (two men).

Stimuli

For all experiments, stimuli consisted of an equal number of grayscale images from each of four different stimulus categories.

Experimental Procedures

Each subject was run on (1) two or more functional localiser scans, and (2) two or more scans from a given experiment. Each fMRI scan lasted a total

duration of 330 seconds and consisted of four blocks of four consecutive 16-second stimulus epochs (one epoch for each stimulus condition) with a 16-second fixation baseline period occurring before each block (e.g. Fig. 1B). Across the four blocks, each stimulus condition appeared once in each serial position within a block. Images from the relevant stimulus condition were serially presented in random sequence during each epoch for subject viewing. During fixation baseline periods, subjects maintained fixation on a central fixation point.

In Experiments 1 and 2, images were centrally presented at a rate of one image every 667msec (stimulus duration = 450msec, interstimulus interval = 217msec) with a small spatial offset (10% of the image width) that alternated between a top-right and bottom-left position. In Experiments 3 and 4, images were centrally presented with no spatial offset at a rate of one image every 800msec (stimulus duration = 400msec, interstimulus interval = 400msec).

Subjects either performed a passive viewing or a one-back matching task in each experiment. For the passive viewing task, subjects were simply instructed to attentively view the sequence of images. In the one-back task, subjects were instructed to press a button whenever they saw two identical pictures in a row. Typically, one or two repetitions occurred in each epoch.

MRI Scanning Procedures

Scanning was done on a 1.5 T scanner at the MGH-NMR Center in Charlestown, MA, using a bilateral quadrature surface coil which provided a high signal to noise ratio in posterior brain regions. Standard echoplanar imaging procedures were used (TR = 2sec, TE = 70msec, flip angle = 90°, 180° offset = 25msec, 165 images/slice). Twelve 6mm or 7mm thick near-coronal slices (parallel to brainstem) covered the entire occipital lobe as well as the posterior portion of the temporal lobes, including the FFA. A bite bar was used to minimise head motion. In all other respects, imaging procedures were identical to those reported by Kanwisher et al. (1997).

fMRI Data Analysis

For each experiment, each subject's FFA was identified on separate functional localiser scans. The FFA was defined as the set of all contiguous voxels in the fusiform gyrus that showed significant differences in activity ($P < .0001$ uncorrected, incorporating a 6-second delay to compensate for haemodynamic lag) on a Kolmogorov-Smirnov test comparing front-view human faces to either houses or familiar objects. Because the localiser data were used to independently define the FFA region of interest for experimental fMRI scans, no correction for multiple voxel-wise comparisons was made.

Each subject's predefined FFA was then used to extract the time course of the magnetic resonance (MR) signal intensity during the experimental fMRI scans (averaging over all voxels identified in the localiser scan). After assuming a 6-second delay in haemodynamic response, FFA responses were then measured in terms of the average percentage change in MR signal during each stimulus epoch compared to fixation as a baseline.

A repeated-measures analysis of variance was performed by pooling the mean FFA response for each subject and stimulus condition with task (i.e. passive or one-back) as a between-subject variable. Planned comparisons were performed between stimulus conditions based on theoretical questions of interest.

Behavioural Data Analysis

The percentage of correct responses on the one-back matching task was calculated for each subject, experimental scan, and stimulus condition in terms of the percentage of hits minus false alarms. Table 1 shows a summary of the mean performance across subjects for Experiments 1–4. Discrimination performance was usually excellent (above 90%) except for stimuli that were perceived as highly visually similar, such as the schematic faces in Experiment 1, the capped faces and eyes alone stimuli in Experiment 3, and the different head views in Experiment 4. Because our main purpose was to study the response properties of the FFA and because we never found an effect of task on FFA responses, the behavioural data were not analysed further.

Table 1. *Summary of Behavioural Performance on one-back Matching Task: Percentage of Hits – False Alarms*

		Stimulus Condition			
Experiment 1	*n*	*Human faces*	*Cat faces*	*Schematic faces*	*Objects*
Mean	4	92	90	72	95
SE		7.8	1.6	4.8	3.0
Experiment 2	*n*	*Human faces*	*Upright cartoon faces*	*Inverted cartoon faces*	*Objects*
Mean	5	94	95	94	98
SE		3.4	3.8	4.0	1.5
Experiment 3	*n*	*Entire faces*	*Faces with no eyes*	*Eyes alone*	*Houses*
Mean	5	78	76	54	79
SE		5.8	6.8	10.9	7.4
Experiment 4	*n*	*Front*	*Profile*	*Cheek*	*Back*
Mean	2	87	75	61	70
SE		3.5	16.7	13.9	4.5

EXPERIMENT 1: HUMAN FACES, CAT FACES, SCHEMATIC FACES, AND OBJECTS

This experiment tested whether the FFA responds selectively to human faces or whether it generalises to other faces such as cat or schematic faces. Examples of the stimuli are shown in Fig. 1A and Appendix A.

Although cat faces possess animate and expressive qualities and are readily perceived as faces, they differ greatly from human faces in terms of their low-level features. For example, cat faces have highly textured fur and are often distinguished based on the patterns of colours in the fur itself. Regarding facial features, cats do not have lips or prominent eyebrows, and their ears, noses and mouths do not resemble those of human faces. The overall shape and aspect ratio of cat faces also differ from human faces. Thus, cat faces allowed us to test whether the FFA is specifically tuned to human facial features or to more general aspects of a face.

The use of cat face stimuli also allowed us to test whether the FFA is involved in expert recognition rather than face perception per se. Some researchers have suggested that the FFA response to human faces may instead reflect an effect of acquired expertise at identifying exemplars within a particular category (e.g. Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1997b). If such a view were correct, one would predict a much greater FFA response to human than cat faces given that people are far more experienced at recognising human faces. However, single-unit studies in monkeys have shown that most face cells generalise equally well across monkey and human faces (e.g. Perrett et al., 1982). It was therefore conceivable that the FFA might also generalise to the faces of other species.

Schematic faces allowed us to test whether a basic facial configuration would be sufficient to activate the FFA. Even newborn infants appear to differentiate between simple schematic faces and scrambled faces, as evidenced by their preference to track a moving schematic face across a greater distance (Goren, Sarty, & Wu, 1975; Johnson, Dziurawiec, Ellis, & Morton, 1991). These results suggest that humans may have an innate preference for simple face-like visual patterns. However, single-unit recordings in monkeys have revealed that face cells show either no response or a weak
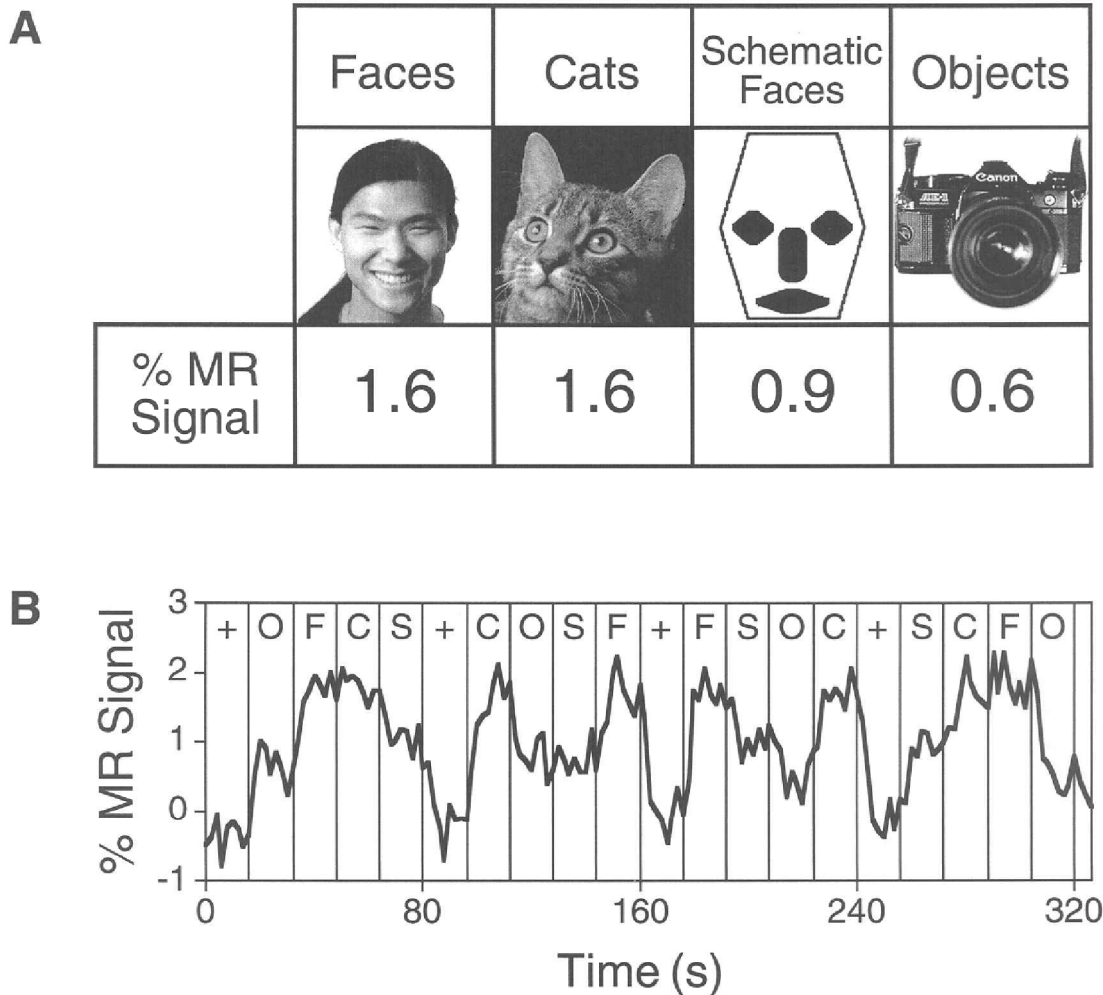
A

	Faces	Cats	Schematic Faces	Objects
% MR Signal	1.6	1.6	0.9	0.6

B

Fig. 1. *1A: Overall FFA response (expressed in percentage MR signal change from fixation baseline) to human faces, cat faces, schematic faces, and objects, averaged across eight subjects in Experiment 1. Example stimuli are shown. 1B: Time course showing mean FFA activity while subjects viewed sequences of human faces (F), cat faces (C), schematic faces (S), and objects (O). Note that FFA responses are delayed by approximately 6 seconds due to haemodynamic delay, and that responses are strongest for human and cat faces in each of the four blocks.*

response to simple line drawings of a monkey face (Perrett et al., 1982).

Unlike the schematic faces used in the above studies, our schematic faces were defined by the arrangement of nonfacial features consisting of simple geometric shapes. This allowed us to investigate the extent to which facial configuration alone, in absence of appropriate facial features, would activate the FFA. Front-view human faces and familiar (inanimate and animate) objects served as optimal and nonoptimal reference conditions respectively.

Method

Subjects

Eight subjects (five women) participated in Experiment 1; four performed the passive viewing task and four performed the one-back matching task.

Stimuli

The four stimulus conditions consisted of human faces, cat faces, schematic faces, and objects. In each condition, 35 different images were used. Figure 1A shows an example of each of these stimuli. Additional examples can be seen in Appendix A. Front-view human faces consisted of digitised college ID photos of undergraduate men and women. Cat faces were taken from a CD-ROM collection of cats and kittens (Corel professional photos). Schematic faces were created by using different rudimentary geometric shapes to compose the eyes, nose, mouth, and external contour of each face. Objects included both easily recognisable inanimate objects (e.g. camera, iron, tape cassette) and animate objects (e.g. cow, horse, caterpillar). All images were cropped and presented within a square window that subtended approximately 12° × 12° of visual angle.

Results and Discussion

Figure 1A shows the overall mean FFA response during each of the four stimulus conditions (expressed in percentage MR signal change from fixation baseline). The FFA responded most strongly to human faces and cat faces, much more weakly to schematic faces, and most weakly to objects. The consistency of these response differences can be seen in the average time course of FFA activity plotted in Fig. 1B. In each of the four blocks, the response to human and cat faces was always greater than the response to schematic faces or objects.

An ANOVA across subjects revealed a highly significant difference among stimulus conditions [$F(3,18) = 54.9$, $P < .00001$]. There was no significant effect of task [$F(1,6) < 1$] or interaction between task and stimulus [$F(3,18) < 1$]. (A summary of mean FFA responses during passive viewing versus one-back matching for all experiments can be found in Table 2.) Planned comparisons revealed significantly greater FFA activity for schematic faces than objects [$t(7) = 3.66$, $P < .01$] and for cat faces than schematic faces [$t(7) = 4.9$, $P < .002$]. However, no significant difference was found between human and cat faces [$t(7) = 1.0$, $P = .35$].

The fact that the FFA responded as strongly to cat faces as human faces is quite striking, given that the FFA was localised on independent scans as the region in the fusiform gyrus that responded significantly more to *human* faces than to objects or houses. These results clearly indicate that the FFA can generalise to faces with very different low-level features, even those of another species. Given that our subjects were far more experienced at discriminating and identifying human faces than cat faces, it seems unlikely that the response of the FFA can be adequately explained in terms of visual expertise alone.

The much weaker response to schematic faces indicates that the presence of facial configuration alone, in absence of appropriate facial features, is not sufficient for fully activating the FFA. The weakness of this response does not seem attributable to the impoverished nature of the schematic faces given that simplistic two-tone Mooney faces have been shown to strongly activate the FFA (Kanwisher, Tong, & Nakayama, 1998), whereas complex objects do not. However, our results do suggest that facial configuration alone may lead to a partial or weak FFA response, as evidenced by the slightly greater response to schematic faces than familiar objects. The present results agree with the finding that face cells in monkeys generally show little or no response to simple schematic faces (Perrett et al., 1982).

EXPERIMENT 2: HUMAN FACES, UPRIGHT CARTOON FACES, INVERTED CARTOON FACES, AND OBJECTS

Experiment 2 tested whether the FFA would generalise to familiar upright and inverted cartoon faces that were loosely based on either animals (e.g. Mickey Mouse) or fictitious characters (e.g. Ernie and Bert; see Fig. 2A and Appendix B for examples of the stimuli). Unlike realistic faces, cartoon faces are drawn in a highly schematic and exaggerated fashion. The facial features of cartoon faces often

Table 2. *Summary of Mean FFA Response as a Function of Stimulus Condition and Task for All Experiments*

		Stimulus Condition			
Experiment 1	*n*	*Human faces*	*Cat faces*	*Schematic faces*	*Objects*
Passive	4	1.8	1.8	1.0	0.7
1-back	4	1.5	1.4	0.8	0.6
Experiment 2	*n*	*Human faces*	*Upright cartoon faces*	*Inverted cartoon faces*	*Objects*
Passive	4	1.7	1.8	1.4	0.7
1-back	5	1.8	1.7	1.4	0.8
Experiment 3	*n*	*Entire faces*	*Faces with no eyes*	*Eyes alone*	*Houses*
Passive	3	2.2	1.9	1.4	1.0
1-back	5	1.6	1.5	1.2	0.5
Experiment 4	*n*	*Front*	*Profile*	*Cheek*	*Back*
Passive	4	1.8	1.8	1.3	0.8
1-back	2	1.8	1.9	1.2	1.2

deviate from normal faces in colour, size, shape, and placement. Despite this fact, cartoons are readily perceived as animate faces, perhaps much more so than the schematic faces in Experiment 1.

There is some neuropsychological evidence suggesting that human faces and cartoon faces access a common face recognition system. Moscovitch, Winocur, and Behrmann (1997) reported an object agnosic patient, CK, who showed spared recognition for human faces as well as familiar cartoon faces. CK could identify both human and cartoon faces as well as normal controls, but was much more severely impaired than normals when the same faces were shown upside-down. Moscovitch et al. suggested that CK had an intact face recognition system that could only operate on upright faces, but an impaired object recognition system that would normally be used for identifying objects and also contribute to the identification of inverted faces. Consistent with this distinction between upright and inverted face processing, Farah, Wilson, Drain, and Tanaka (1995) reported a prosopagnosic patient who showed the opposite deficit of impaired recognition for upright faces but normal recognition for inverted faces.

This experiment used the same cartoon faces that were tested on patient CK to see if upright cartoon faces would activate the FFA as strongly as upright human faces. We also presented the cartoon faces upside-down to see if this would result in weaker FFA activity and, if so, whether FFA activity would drop to the level found for objects. Given the neuropsychological and cognitive evidence that face-specific processing is severely disrupted for inverted faces (Farah et al., 1995; Moscovitch et al., 1997; Tanaka & Farah, 1993; Tong & Nakayama, in press; Yin, 1971; Young, Hellawell, & Hay, 1987), one might expect that inverted faces should fail to access face-specific mechanisms. However, single-unit recordings in monkeys have revealed mixed results. Whereas most face cells in the lateral inferotemporal cortex respond more to upright than inverted faces (Tanaka, Saito, Fukada, & Moriya, 1991), cells in the superior temporal sulcus respond equally strongly to both types of faces (Perrett et al., 1982, 1985).

In previous fMRI studies of human observers, we have found that FFA responses are only slightly stronger for upright than inverted grayscale human faces, even though face discrimination performance

is much poorer for the inverted faces (Kanwisher et al., 1998). By contrast, two-tone Mooney faces, which are difficult to perceive as faces at all after inversion, were found to yield a larger and more consistent FFA inversion effect.

Given that the subjects in the present experiment readily saw the inverted cartoons as faces, we predicted that the FFA inversion effect for cartoon faces would be rather small. In fact, even after inversion the cartoons often appeared easy to discriminate and identify because they contained highly distinctive features (e.g. the ears of Mickey Mouse or Bugs Bunny). We were therefore curious whether distinctive cartoon faces would still yield an FFA inversion effect.

Method

Subjects

Nine subjects (seven women) participated in Experiment 2; four performed the passive viewing task and five performed the one-back matching task.

Stimuli

Stimulus conditions consisted of upright cartoon faces (e.g. Mickey Mouse, Big Bird, Ernie, and Bert), inverted cartoon faces (i.e. the same cartoons shown upside-down), human faces (college ID photos), and objects. In each condition, 34 different images were used. All images were cropped and presented within a square window that subtended approximately 12° × 12° of visual angle. See Fig. 2A and Appendix B for examples.

Results and Discussion

FFA responses were strongest for upright human faces and upright cartoon faces, slightly weaker for inverted cartoon faces, and weakest for objects (see Fig. 2A). The consistency of these activation differences, including the cartoon inversion effect, could be seen in the average fMRI time course (see Fig. 2B) as well as in the individual subject data.

An ANOVA revealed a highly significant difference among stimulus conditions [$F(3,21) = 42.9$, $P < .00001$] but no significant effect of task or interaction between task and stimulus [$F(1,7) < 1$] or interaction between task and stimulus [$F(3,21) < 1$]. Planned comparisons revealed negligible differences in FFA response to upright cartoon faces and human faces [$t(8) = 0.56$, $P = .59$]. This indicates that the FFA response generalises equally well to upright cartoon faces as to human or cat faces (Experiment 1) despite their very different low-level image properties. These results are consistent with the preserved face recognition found in object agnosic patient CK, who could recognise both upright human and cartoon faces equally well (Moscovitch et al., 1997).

We found a small but significant decrease in FFA activity for inverted cartoon faces relative to upright cartoon faces [$t(8) = -6.74$, $P < .0001$] despite negligible differences in behavioural performance on the one-back matching task (inverted cartoons = 94% correct vs. upright cartoons = 95% correct). Compared to the small FFA inversion effect (upright cartoons − inverted cartoons = 1.7% − 1.4% = 0.3%), the FFA response to inverted cartoon faces was much greater than the response found for objects [1.4% vs. 0.7% respectively, $t(8) = 4.70$, $P < .002$], in fact almost twice in magnitude. The weak FFA inversion effect found here contrasts with the severe recognition deficit that CK showed for inverted cartoon faces (Moscovitch et al., 1997), but agrees well with the rather weak FFA inversion effects previously found for grayscale human faces (Kanwisher et al., 1998). Our results suggest that inverted faces can access face-specific mechanisms such as the FFA to a considerable extent, though not quite as effectively as upright faces.

EXPERIMENT 3: ENTIRE FACES, FACES WITHOUT EYES, EYES ALONE, AND HOUSES

Experiment 3 investigated the extent to which the FFA response could be attributed to the presence of the eyes. This was done by comparing the FFA response for entire human faces, faces with eyes occluded, eyes alone, and houses (see Fig. 3A and Appendix C for examples of the stimuli).
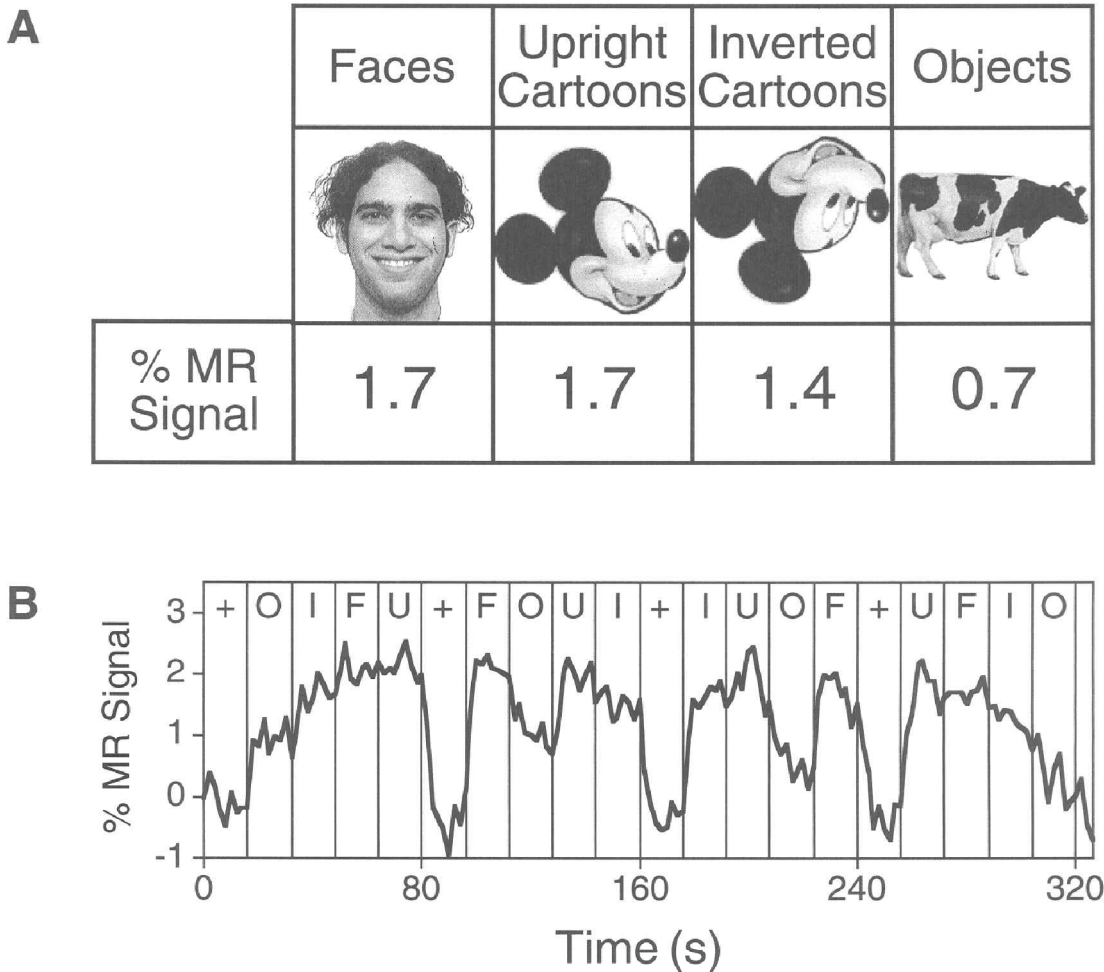
A

	Faces	Upright Cartoons	Inverted Cartoons	Objects
% MR Signal	1.7	1.7	1.4	0.7

B

Fig. 2. *Example stimuli, overall FFA response (panel A) and time course of FFA activity (panel B) averaged across nine subjects in Experiment 2 during visual presentation of human faces (F), upright cartoon faces (U), inverted cartoon faces (I), and objects (O). Note that FFA responses are strongest for upright human and cartoon faces, somewhat weaker for inverted cartoon faces, and weakest for objects in all four blocks.*

One possible explanation of the generalised FFA response found across human, cat, and cartoon faces is that it reflects a more basic response to a salient facial feature such as the eyes. Perrett et al. (1985) found that for many "face cells" in superior temporal sulcus, selective responses for a particular face view (front or profile) could be better attributed to the gaze direction of the eyes. Furthermore, lesions of the superior temporal sulcus have been found to impair the discrimination of gaze direction but have little effect on face recognition performance, suggesting that this region of monkey cortex may be specifically involved in gaze perception (Heywood & Cowey, 1992). ERP studies in humans have also revealed an N170 potential over the posterior temporal scalp that is greater for the eyes than for the nose, mouth, or entire face (Bentin, Allison, Puce, Perez, & McCarthy, 1996). Such an ERP may reflect an underlying gaze detection mechanism. However, such a mechanism has yet to be localised in the human brain. By using fMRI, we were able to test whether the FFA is

selectively involved in gaze perception or involved in more global aspects of face perception.

Method

Subjects
Eight subjects (four women) participated in Experiment 3; three performed the passive viewing task and five performed the one-back matching task.

Stimuli
Front-view digital images of 38 men wearing a black ski hat were taken under controlled laboratory conditions. Stimuli consisted of 19 images of entire faces, a different set of 19 faces from which the eyes were digitally removed (and used for the eyes alone stimuli), 19 images of eyes alone, and 19 houses. All images were cropped and presented within a square window that subtended approximately 16° × 16° of visual angle, except for the eyes alone stimuli which subtended approximately 6° × 1.8°. See Fig. 3A and Appendix C for examples.

Results and Discussion

FFA responses were strongest for entire faces and faces without eyes, weaker for eyes alone, and weakest for houses (see Fig. 3A). The consistency of these activation differences can be seen in the average fMRI time course for each of the four blocks (see Fig. 3B).

An ANOVA revealed a highly significant difference among stimulus conditions [$F(3,18) = 28.5$, $P < .00001$], but no significant effect of task [$F(1,6) = 1.0$, $P = .30$] or interaction between task and stimulus [$F(3,18) < 1$]. Planned comparisons revealed a negligible difference in FFA response to entire faces and faces without eyes [$t(7) = 1.3$, $P = .23$], but a significantly greater response to faces without eyes than eyes alone [$t(7) = 3.30$, $P < .02$]. These results indicate that the eyes are neither necessary nor sufficient to fully activate the FFA. It further suggests that the FFA is not selectively involved in gaze detection (cf. Bentin et al., 1996; Heywood & Cowey, 1992), but instead is likely to be involved in global aspects of face perception. It

should also be noted that we found no other brain regions in our posterior slices that showed consistently greater responses to the presence of the eyes based on comparisons between eyes versus houses or entire faces versus faces without eyes.

The FFA response to eyes alone was much greater than the response to houses [$t(7) = 6.3$, $P < .0005$], even though the house stimuli were much larger and more visually complex. Apparently, a small isolated facial feature such as the eyes can still evoke a fairly large FFA response. However, when this feature was added to a face without eyes to form a complete face, no increase in FFA response was observed. This suggests that the FFA does not combine information from various facial features in a strictly linear fashion (see also Perrett et al., 1982, 1985).

EXPERIMENT 4: FRONT, PROFILE, CHEEK, AND BACK OF HEAD VIEWS

Experiment 4 investigated whether the FFA response is specific to faces or whether it reflects a more general response to a broader class of stimuli such as different views of heads. We tested four different head views: front view (0° rotation), profile view (90°), cheek view (135°), and back view (180°). Examples of the stimuli are shown in Fig. 4A and Appendix D.

A possible explanation of the generalised FFA response found across cat, cartoon, and human faces is that the FFA is more generally tuned to heads. For example, Sinha and Poggio (1996) have suggested that head recognition rather than face recognition is often used to identify individuals. By contrast, Perrett et al. (1991) have suggested that the reason why most "face cells" in the macaque superior temporal sulcus show specificity for various head views (including back of head and cheek views) is that these neurons are involved in coding the direction of another person's attention rather than facial identity.

We reasoned that if the FFA is generally tuned to heads, then one would predict equally elevated responses across all head views. However, if the
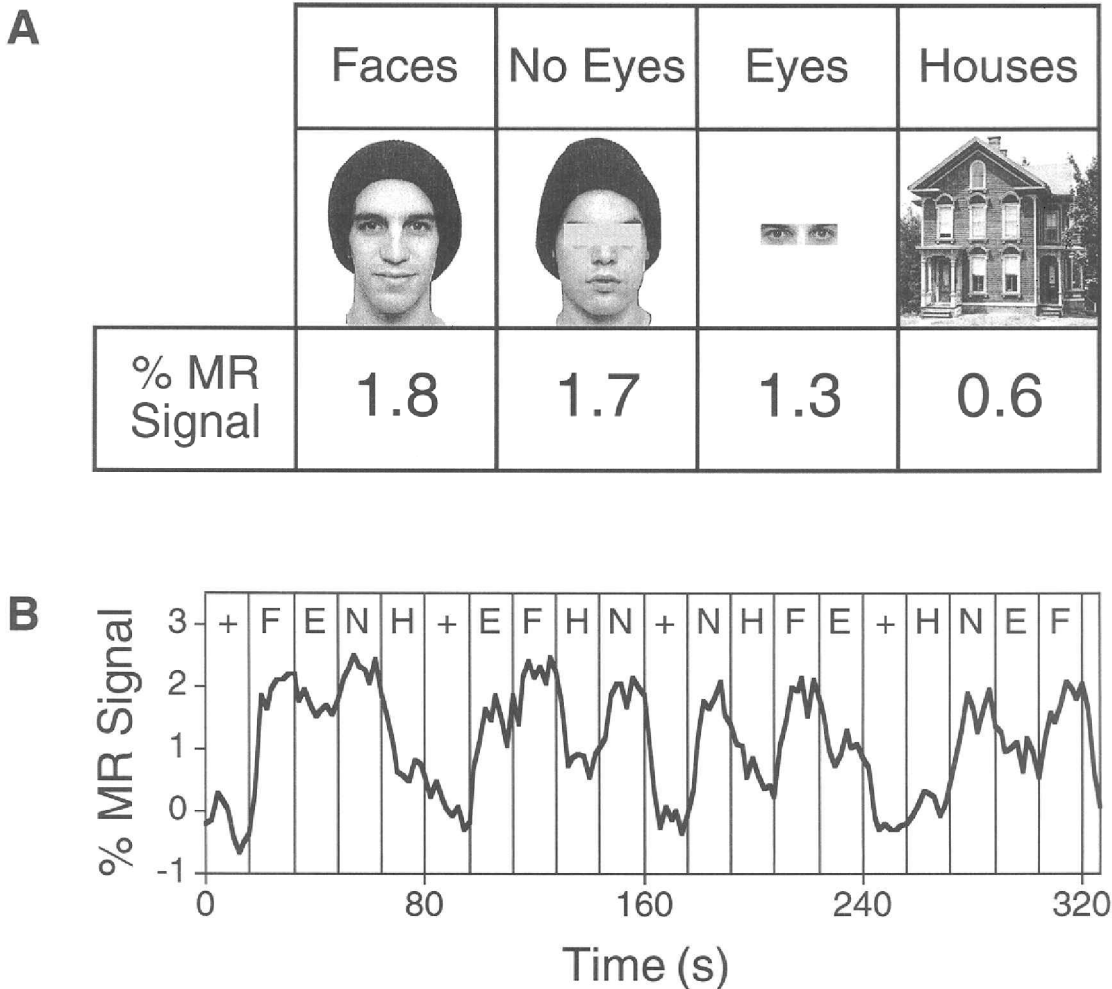
A

	Faces	No Eyes	Eyes	Houses
% MR Signal	1.8	1.7	1.3	0.6

B

Fig. 3. *Example stimuli, overall FFA response (panel A) and time course of FFA activity (panel B) averaged across eight subjects in Experiment 3 during visual presentation of entire human faces (F), faces with no eyes (N), eyes alone (E), and houses (H). Note that FFA responses are strongest for entire faces and faces with no eyes, weaker for eyes alone, and weakest for houses in all four blocks.*

FFA is more specifically tuned to faces, one would predict equally strong activity for front and profile views in which the internal features of the face remain visible, but declining activity thereafter as the face becomes progressively hidden from view.

Method

Subjects
Five subjects (two women) participated in Experiment 4. Three subjects performed only the passive viewing task, one subject performed only the one-back task, and one subject performed both tasks on different scans in the same session.

Stimuli
Images of 13 different men were taken from 4 different viewpoints which served as the 4 different stimulus conditions: front view (0° rotation), profile view (90°), cheek view (135°), and back view (180°). All images were cropped and presented within a

square window that subtended approximately 12° × 12° of visual angle. See Fig. 4A and Appendix D for examples of the stimuli.

Results and Discussion

The FFA responded most strongly to front and profile views, more weakly to cheek views and most weakly to back of head views (see Fig. 4A). The consistency of these activation differences can be seen in the average fMRI time course for each of the four blocks (see Fig. 4B).

An ANOVA revealed a highly significant difference among stimulus conditions [$F(3,12) = 12.9$, $P < .0005$], but no significant effect of task [$F(1,4) < 1$] or interaction between task and stimulus [$F(3,12) < 1$]. Planned comparisons revealed negligible FFA response differences between front and profile views [$t(5) = -0.08, P = .94$], but significantly greater activity for profile than cheek views [$t(5) = 4.01, P < .02$] and for cheek views than back of head views [$t(5) = 2.86, P < .05$]. Although an object control condition was not included in this experiment, the magnitude of FFA response to backs of heads appeared comparable to or only slightly greater than the response typically found for non-face objects.

FFA activity declined as the head rotated beyond profile view and the internal features of the face became progressively more occluded. This strongly suggests that the FFA responds specifically to faces and not more generally to heads. These results also rule out the possibility that the generalised FFA response found for cat and cartoon faces reflects a response to heads or animate forms in general. All subjects clearly recognised that they were viewing the back of human heads, and some subjects who were familiar with the persons shown even reported that they could identify people from the back of their heads. Despite the fact that such views were readily perceived as animate and were sometimes even recognised as belonging to a certain individual, the FFA failed to show much response to these views.

GENERAL DISCUSSION

The experiments reported here show that the FFA response generalises equally well across cat faces, cartoon faces, front-view human faces (with or without eyes), and profile-view human faces, even though the low-level visual features that compose these stimuli differ greatly. Such similar responses to very heterogeneous stimuli cannot be explained in terms of a more specific response to a salient facial feature such as the eyes or a more general response to heads or animate objects. Instead, our results strongly suggest that the FFA is optimally tuned to the broad stimulus category of faces.

Whereas cat, cartoon, and human faces evoked the strongest FFA responses, stimuli that conveyed only certain perceptual characteristics of a face (e.g. simple schematic faces, inverted cartoon faces, eyes alone, and cheek views of heads) evoked intermediate responses. By contrast, nonface stimuli such as objects, houses, and backs of heads consistently yielded the weakest responses. These results have a number of implications regarding the functional role of the FFA and the underlying processes that it may be performing. We address these issues below.

Is the FFA Selectively Involved in Gaze Perception?

There is some evidence that superior temporal sulcus lesions in monkeys may selectively impair the ability to perceive the direction of another's gaze without affecting face recognition performance (Heywood & Cowey, 1992). ERP recordings in humans have also revealed an N170 potential over the lateral occipital regions that is greater for the eyes than for the nose, mouth, or entire face (Bentin et al., 1996). These studies suggest that there may exist neural regions dedicated to the more specific task of gaze perception.

However, our results in Experiment 3 revealed no evidence of neural regions that respond selectively to gaze. Outside of the FFA, we failed to find preferential responses to the eyes. Note, however, that our brain slices extended from the occipital pole to the posterior temporal lobes; therefore it

A

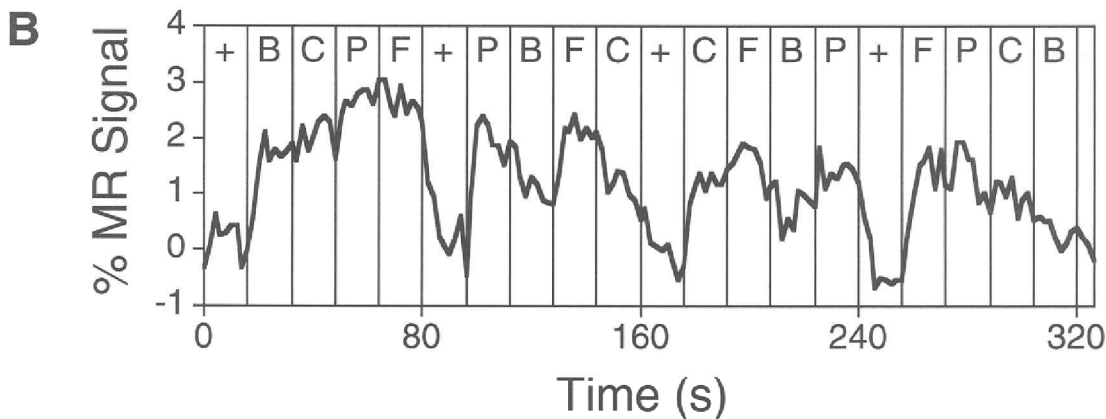| % MR Signal	Front	Profile	Cheek	Back
	1.8	1.8	1.3	0.9

B

Fig. 4. *Example stimuli, overall FFA response (panel A) and time course of FFA activity (panel B) averaged across six subjects in Experiment 4 during visual presentation of front (F), profile (P), cheek (C), and back (B) views of human heads. Note that FFA responses are strongest for front and profile views, weaker for cheek views, and weakest for back views in all four blocks.*

remains possible that one might find gaze-selective responses in more anterior regions.

Within the FFA itself, fMRI responses were equally strong for entire faces and faces without eyes, but significantly weaker for eyes shown alone. Thus, the eyes are neither necessary nor sufficient for evoking an optimal FFA response. This suggests that FFA is not *selectively* involved in gaze detection but instead is involved in more global aspects of face perception.

Does FFA Activity Reflect the Processing of Facial Features, Facial Configuration, or Both?

Although the processing of individual facial features contributes to face recognition, cognitive studies have generally emphasised the importance of configural or holistic processing for effective face recognition (Rhodes, Brennan, & Carey, 1987; Tanaka & Farah, 1993; Young et al., 1987). How-

ever, recent investigation of object agnosic patient CK, who appears to have a selectively spared face recognition system, suggests that both configural and featural information are encoded within a unitary system (Moscovitch et al., 1997). Our study likewise suggests that both global configuration and local facial features serve to activate the FFA.

In Experiment 1, we found that schematic faces defined by simple geometric shapes activated the FFA only slightly more than nonface objects. The weak response to schematic faces indicates that the presence of facial configuration alone, in the absence of appropriate facial features, is insufficient for evoking an optimal FFA response but can lead to a partial response.

In Experiment 2, subjects viewed upright and inverted cartoon faces that provided the same low-level feature information. However, FFA responses were significantly weaker for the inverted cartoon faces (see also Kanwisher et al., 1998), most probably because configural processing was disrupted for these faces. It is well known that face inversion severely disrupts configural or holistic processing, but generally has little effect on piecemeal processing of individual facial features (Tanaka & Farah, 1993; Tanaka & Sengco, 1997; Young et al., 1987).

These results suggest that neither configural information alone nor feature information alone is sufficient for evoking an optimal FFA response. Instead, optimal FFA activation appears to require both appropriate facial features and appropriate facial configuration. These response properties may reflect the fact that different FFA neurons are activated by different aspects of the face. Single-unit recordings in monkeys have revealed that face cells may be tuned to a particular facial feature (e.g. eyes, nose, mouth, chin), combination of features (e.g. mouth and chin) or the overall configuration of a face (Perrett et al., 1982; Yamane et al., 1988). If FFA neurons possess similar variations in tuning, one would predict that the FFA population would be most fully activated by a complete set of facial features presented in a proper configuration. By contrast, one would predict that fewer FFA neurons would respond to a face with inappropriate features (e.g. schematic faces), inappropriate con-figuration (e.g. inverted faces), or to an isolated facial feature (e.g. eyes alone), as was found here.

Why Does the FFA Respond to Inverted Cartoon Faces?

Although we found a significant inversion effect for cartoon faces, the response to inverted cartoons remained almost twice as strong as the response to nonface objects. Perhaps one reason why inverted cartoon faces and inverted grayscale human faces (Kanwisher et al., 1998) evoke considerable FFA activity is that many facial features such as the eyes remain highly salient even after configural processing is disrupted by inversion. As Experiment 3 revealed, the eyes alone can evoke a fairly large FFA response.

Another possible explanation, which does not exclude the preceding one, is that the FFA is largely activated by successful acts of face *perception* rather than successful acts of face *recognition*. Inversion generally leads to rather small decreases in FFA activity, regardless of whether recognition performance is severely impaired, as in the case of grayscale human faces (Kanwisher et al., 1998), or unimpaired, as was found for cartoon faces. Even when these faces are inverted, they are still readily perceived as faces. By contrast, two-tone Mooney faces, which are difficult to perceive as faces once inverted, yield stronger FFA inversion effects (Kanwisher et al., 1998). This suggests that the FFA may be primarily activated by the perception of a face.

Supporting this view, Tong, Nakayama, Vaughan, and Kanwisher (1998) recently found that when subjects continuously viewed a bistable face/house display under conditions of binocular rivalry, the FFA showed a sharp increase in activity whenever subjects reported a perceptual switch from house to face. By contrast, the FFA showed a sharp decrease in activity during perceptual switches from face to house. FFA activity may therefore be tightly related to the phenomenal experience of seeing a face, even when it is the same face that is repeatedly perceived.

Might the FFA be Primarily Involved in Subordinate-level Categorisation or Expert Recognition?

Are there alternative accounts that might explain the apparent face-specificity of the FFA? Gauthier, Anderson, Tarr, Skudlarski, and Gore (1997a) have suggested that the FFA may reflect subordinate-level categorisation within a homogenous visual category rather than face-specific processing. However, such an account seems unlikely given that Kanwisher et al. (1997) previously showed that FFA responses were almost four times greater when subjects discriminated between individual faces than between individual hands. Here, we found that FFA responses were strongest for cat faces, upright cartoons, and human faces, intermediate for schematic faces, inverted cartoons, eyes alone, and cheek views, and weakest for houses and backs of heads, irrespective of whether subjects performed subordinate-level discriminations in the one-back matching task or a simple passive viewing task (see Table 2 for summary). These results suggest that the FFA response is quite automatic and far more dependent upon stimulus attributes than upon the difficulty of the visual discriminations involved (see Table 1) or the task at hand.

More recently, Gauthier et al. (1997b) have suggested that the FFA may instead reflect subordinate-level categorisation within a visual category for which one has acquired visual expertise. We think that such a modified account is also unlikely given the very strong FFA responses found here for cat faces. Although all of our subjects were far more experienced at recognising and individuating human faces than cat faces, the FFA showed equally strong activations to both face types. An expertise account would also have difficulty explaining why letters, which represent another highly overlearned visual category, activate different brain regions than faces (Puce et al., 1996). These findings suggest that the FFA response cannot be adequately explained in terms of the degree of visual expertise that one has for a particular category. Although some other modified subordinate-level classification account may still be viable, we

believe that the simplest account requiring the fewest ad hoc assumptions is that the FFA is primarily engaged in face perception.

Possible Roles of the FFA in the Detection, Perception, and Recognition of Faces

At present, there is little evidence to suggest that the FFA is involved in recognition memory for individual faces. Neuroimaging studies have shown that the FFA responds equally well to familiar versus novel faces, suggesting that the FFA might not be the site of memory storage for faces (Clark, Maisog, & Haxby, 1998; Haxby et al., 1996). These studies do not preclude the possibility that the FFA is involved in recognition memory for faces, but they do not support such involvement.

A plausible alternative is that the FFA plays a role in the *detection* of faces. By face detection, we mean the ability to discriminate between faces and nonfaces. The present study revealed that the FFA is optimally activated by cat, cartoon, and human faces but is weakly activated by nonface stimuli. This suggests that the FFA may not be specialised for recognising human faces but instead may serve a more basic function in detecting any type of face. In other studies, we have found that face inversion most severely impairs the FFA response when it impairs the subject's ability to detect the presence of a face (Kanwisher et al., 1998) and that the FFA automatically responds when the same face is repeatedly perceived in a bistable face/house display (Tong et al., 1998). These results suggest that the FFA response may reflect the conscious detection of a face.

Why might people develop a specialised mechanism for face detection? A face detection mechanism could signal the presence of humans or other animals in the immediate environment and thus serve as an alerting mechanism, providing important information for adaptive behaviour and survival. Face detection may also provide the necessary input for subsequent face recognition processes. For many computer vision algorithms, a face must first be located in a scene before face identification

can proceed (e.g. *FaceIt*, 1998; Turk & Pentland, 1991) and, surprisingly, face detection in complex static scenes can sometimes prove as challenging as face identification itself (Leung, Burl, & Perona, 1995).

The FFA may also play a role in *face perception*, perhaps by representing the underlying shape or structure of a face. Inversion of grayscale human faces impairs both the FFA response (Kanwisher et al., 1998) and the perception of facial configuration (Tanaka & Sengco, 1997; Young et al., 1987), even though such inverted images are readily detected as faces. Thus, FFA activity may not only reflect face detection but may also reflect the perception of a face's shape and configuration.

One reason why face detection and face perception may rely on a common mechanism is computational efficiency. A common representational scheme such as a multidimensional "face space" (e.g. Rhodes et al., 1987; Valentine, 1991) could effectively serve both functions. Face detection would involve determining whether a particular item sufficiently resembles the central tendency of faces, whereas face perception would involve the complementary process of describing how a face deviates from this central tendency. Interestingly, when subjects must distinguish between intact and scrambled faces in a speeded-response task, faces that are highly atypical and recognisable are actually more difficult to classify as faces (Valentine, 1991). This suggests that face detection, like face perception, may involve coding a face in terms of its deviation from the central tendency of faces.

According to this framework, the FFA may play an important role in face detection and face perception, and may ultimately serve as a gateway to the higher-level areas where memories for individual faces are stored. FFA activity presumably reflects the activation of face-selective units that signal the presence of a face. These units may also provide a structural code for each face, perhaps by representing the local features and global configuration. Such structural information would provide the necessary input to higher-level recognition areas that match the incoming face input to stored face memories.

Concluding Remarks

We have shown that the FFA responds optimally to a variety of faces (humans, cats, and cartoons) and less strongly to stimuli that convey only certain perceptual characteristics of faces (e.g. schematic faces, inverted cartoons, eyes alone, cheek views of heads). The response profile of the FFA cannot be adequately explained in terms of low-level image properties given that upright and inverted faces with common image properties evoke different responses, and given that cat, cartoon, and human faces with very different image properties evoke comparable responses. Instead, the FFA response profile may be better understood in terms of higher-order representations for the general category of faces. One possibility is that activity in the FFA is tightly linked to the phenomenal experience of detecting or perceiving a face (Tong et al., 1998).

Regarding the functional role of the human fusiform face area, our results suggest that the FFA is not primarily involved in gaze perception, head detection, subordinate-level categorisation, or expert recognition. Instead, the FFA appears to be strongly involved in face detection and face perception, and may play a role in processing the local features and global configuration of faces.

REFERENCES

Allison, T., Ginter, H., McCarthy, G., Nobre, A.C., Puce, A., Luby, M., & Spencer, D.D. (1994). Face recognition in human extrastriate cortex. *Journal of Neurophysiology*, *71*, 821–825.

Bentin, S., Allison, T., Puce, A., Perez, E., & McCarthy, G. (1996). Electrophysiological studies of face perceptions in humans. *Journal of Cognitive Neuroscience*, *8*, 551–565.

Bodamer, J. (1947). Die Prosopagnosie. *Archiv für Psychiatrie und Nervenkrankheiten*, *179*, 6–53.

Clark, V.P., Maisog, J.M., & Haxby, J.V. (1998). fMRI study of face perception and memory using random stimulus sequences. *Journal of Neurophysiology*, *79*, 3257–3265.

FaceIt (1998). [Computer software.] Jersey City, NJ: Visionics Corporation.

Farah, M.J., Wilson, K.D., Drain, H.M., & Tanaka, J.R. (1995). The inverted face inversion effect in prosopagnosia: Evidence for mandatory, face-specific perceptual mechanisms. *Vision Research*, *35*, 2089–2093.

Fried, I., MacDonald, K.A., & Wilson, C.L. (1997). Single neuron activity in human hippocampus and amygdala during recognition of faces and objects. *Neuron, 18*, 753–765.

Gauthier, I., Anderson, A.W., Tarr, M.J., Skudlarski, P., & Gore, J.C. (1997a). Levels of categorization in visual recognition studied using function magnetic resonance imaging. *Current Biology*, *7*, 645–651.

Gauthier, I., Tarr, M.J., Anderson, A.W., Skudlarski, P., & Gore, J.C. (1997b). Expertise training with novel objects can recruit the fusiform face area. *Society for Neuroscience Abstracts*, *23*(2), 2229.

Goren, C.C., Sarty, M., & Wu, P.Y.K. (1975). Visual following and pattern discrimination of face-like stimuli by newborn infants. *Pediatrics*, *56*, 544–549.

Gross, C.G., Roche-Miranda, G.E., Bender, D.B. (1972). Visual properties of neurons in the infero-temporal cortex of the macaque. *Journal of Neurophysiology*, 35, 96–111.

Hasselmo, M.E., Rolls, E.T., & Baylis, G.C. (1989). The role of expression and identity in the face-selective responses of neurons in the temporal visual cortex of the monkey. *Behavioral Brain Research*, *32*, 203–218.

Haxby, J.V., Horwitz, B., Ungerleider, L.G., Maisog, J.M., Pietrini, P., & Grady, C.L. (1994). The functional organization of human extrastriate cortex: A PET-rCBF study of selective attention to faces and locations. *Journal of Neuroscience*, *14*, 6336–6353.

Haxby, J.V., Ungerleider, L.G., Horwitz, B., Maisog, J.M., Rapoport, S.I., & Grady, C.L. (1996). Face encoding and recognition in the human brain. *Proceedings of the National Academy of Sciences USA*, *93*, 922–927.

Heit, G., Smith, M.E., & Halgren, E. (1988). Neural encoding of individual words and faces by the human hippocampus and amygdala. *Nature*, *333*, 773–775.

Heywood, C.A., & Cowey, A. (1992). The role of the 'face-cell' area in the discrimination and recognition of faces by monkeys. *Philosophical Transactions of the Royal Society of London: Biological Sciences*, *335*(1273), 31–37.

Jeffreys, D.A. (1989). A face-responsive potential recorded from the human scalp. *Experimental Brain Research*, *78*, 193–202.

Jeffreys, D.A. (1996). Evoked potential studies of face and object processing. *Visual Cognition*, *3*, 1–38.

Johnson, M.H., Dziurawiec, S., Ellis, H., & Morton. J. (1991). Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition*, *40*, 1–19.

Kanwisher, N., McDermott, J., & Chun, M.M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, *17*(11), 4302–4311.

Kanwisher, N., Tong, F., & Nakayama, K. (1998). The effect of face inversion on the human fusiform face area. *Cognition*, *68*, B1–B11.

Leung, T.K., Burl, M.C., & Perona, P. (1995, June). *Finding faces in cluttered scenes using random graph matching*. International Conference on Computer Vision, Cambridge, MA.

McCarthy, G., Puce, A., Gore, J.C., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. *Journal of Cognitive Neuroscience*, *9*(5), 604–609.

Meadows, J.C. (1974). The anatomical basis of prosopagnosia. *Journal of Neurology, Neurosurgery and Psychiatry*, *37*, 489–501.

Moscovitch, M., Winocur, G., & Behrmann, M. (1997). What is special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. *Journal of Cognitive Neuroscience*, *9*(5), 555–604.

Ojemann J.G., Ojemann, G.A., Lettich, E. (1992). Neuronal activity related to faces and matching in human right nondominant temporal cortex. *Brain*, *115*, 1–13.

Perrett, D.I., Oram, M.W., Harries, M.H., Bevan, R., Hietanen, J.K., Benson, P.J., & Thomas, S. (1991). Viewer-centred and object-centred coding of heads in the macaque temporal cortex. *Experimental Brain Research*, *86*, 159–173.

Perrett, D.I., Rolls, E.T., & Caan, W. (1982). Visual neurones responsive to faces in the monkey temporal cortex. *Experimental Brain Research*, *47*, 329–342.

Perrett, D.I., Smith, P.A.J., Potter, D.D., Mistlin, A.J., Head, A.S., Milner, A.D., & Jeeves, M.A. (1985). Visual cells in the temporal cortex sensitive to face view and gaze direction. *Proceedings of the Royal Society of London: B*, *223*, 293–317.

Puce, A., Allison, T., Asgari, M., Gore, J.C., & McCarthy, G. (1996). Differential sensitivity of human visual cortex to faces, letterstrings, and textures: A functional magnetic resonance imaging study. *Journal of Neuroscience*, *16*, 5205–5215.

Rhodes, G., Brennan, S., & Carey, S. (1987). Identification and ratings of caricatures: Implications for mental representations of faces. *Cognitive Psychology, 19*(4), 473–497.

Sams, M., Hietanen, J.K., Hari, R., Ilmoniemi, R.J., Lounasmaa, O.V. (1997). Face-specific responses from the human inferior occipito-temporal cortex. *Neuroscience, 77,* 49–55.

Sergent, J., Ohta, S., MacDonald, B. (1992) Functional neuroanatomy of face and object processing: a positron emission tomography study. *Brain, 115,* 15–36.

Sinha, P., & Poggio, T. (1996). I think I know that face. *Nature, 384,* 404.

Tanaka, J.W., & Farah, M.J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology, 46A*(2), 225–245.

Tanaka, J.W., & Sengco, J.A. (1997). Features and their configuration in face recognition. *Memory and Cognition, 25*(5), 583–592.

Tanaka, K., Saito, H.-A., Fukada, Y., & Moriya, M. (1991). Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *Journal of Neurophysiology, 66,* 170–189.

Tong, F., & Nakayama, K. (in press). Robust representations for faces: Evidence from visual search. *Journal of Experimental Psychology: Human Perception and Performance.*

Tong, F., Nakayama, K., Vaughan, J.T., & Kanwisher, N. (1998). Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron, 21,* 753–759.

Turk, M. & Pentland, A. (1991). Eigenfaces for recognition. *Journal of Cognitive Neuroscience. 3,* 71–86.

Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Quarterly Journal of Experimental Psychology, 43A,* 161–204.

Yamane, S., Kaji, S., & Kawano, K. (1988). What facial features activate face neurons in the inferotemporal cortex of the monkey. *Experimental Brain Research, 73,* 209–214.

Yin, R.K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology, 81,* 141–145.

Young, A.W., Hellawell, D., & Hay, D.C. (1987). Configural information in face perception. *Perception, 16,* 747–759.
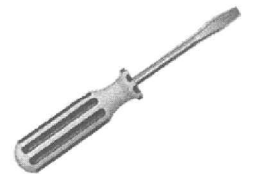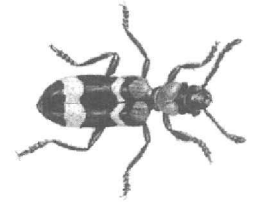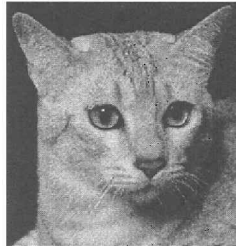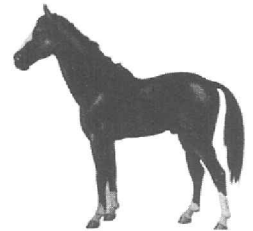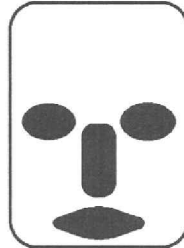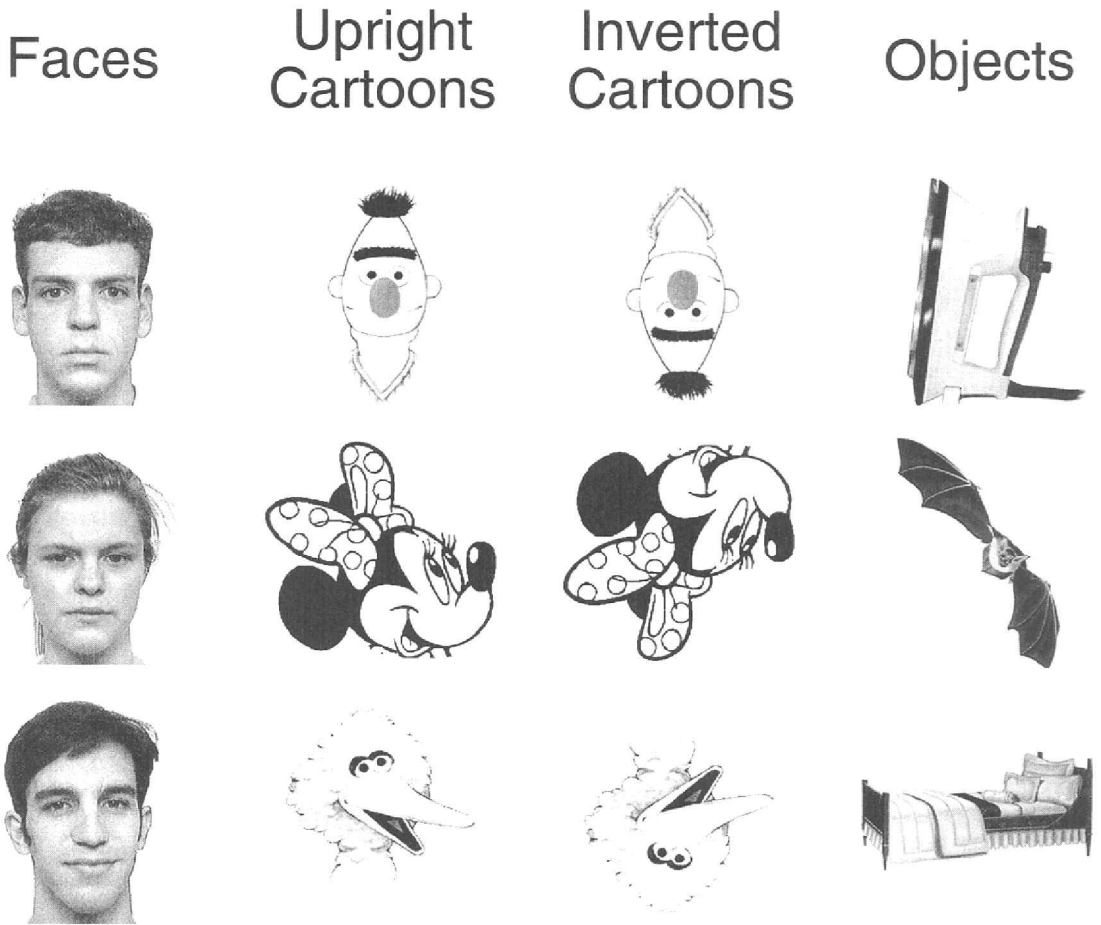
APPENDIX A

Examples of stimuli in Experiment 1.

Faces　　　Cats　　　Schematic Faces　　　Objects

APPENDIX B

Examples of stimuli in Experiment 2.

Faces	Upright Cartoons	Inverted Cartoons	Objects

APPENDIX C

Examples of stimuli in Experiment 3.

Faces No Eyes Eyes Houses
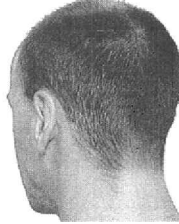
APPENDIX D

Examples of stimuli in Experiment 4.
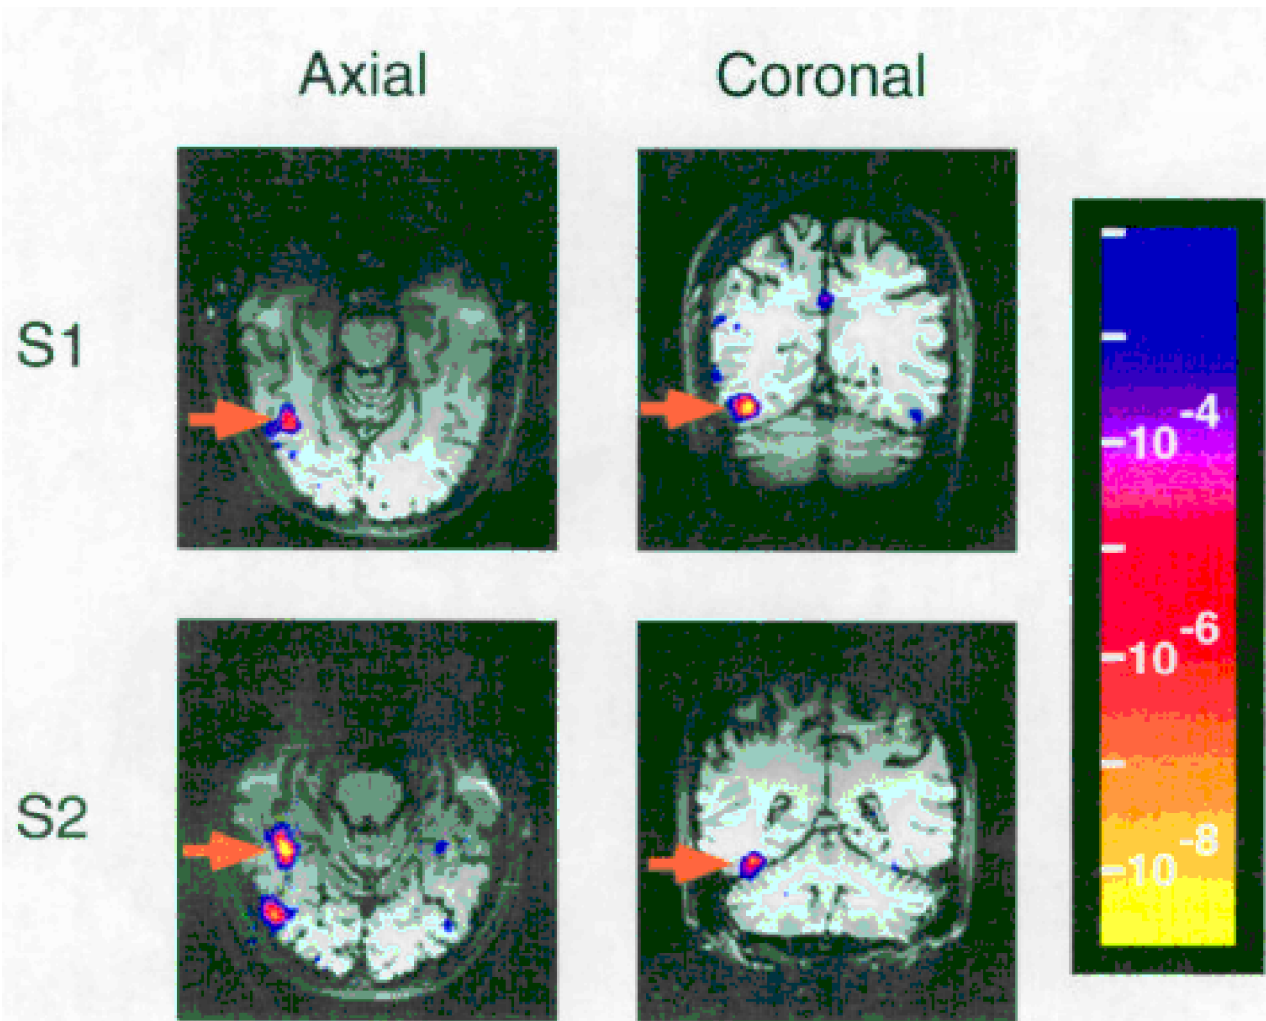

Front   Profile   Cheek   Back
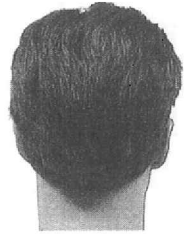
Plate 3 (Tong et al.). *Axial and coronal slices showing the fusiform face area (FFA) in two subjects. Arrows point to the FFA, which included all contiguous voxels in the fusiform gyrus that showed significant differences in activity to faces minus objects or faces minus houses (P < .0001 uncorrected on a Kolmogorov–Smirnov test, see adjacent colour scale) on functional localiser scans. Note that the right hemisphere is shown on the left side in all images and vice versa*.