Heterogeneous Structure in Face-selective Human Occipito-temporal Cortex

Lisa R. Betts* and Hugh R. Wilson

Abstract

■ It is well established that the human visual system contains a distributed network of regions that are involved in processing faces, but our understanding of how faces are represented within these face-sensitive brain areas is incomplete. We used fMRI to investigate whether face-sensitive brain areas are solely tuned for whole faces, or whether they contain heterogeneous populations of neurons tuned to individual components of the face as well as whole faces, as suggested by physiological investigations in nonhuman primates. The middle fusiform gyrus (fusiform face area, or FFA) and the inferior occipital gyrus (occipital face area, or OFA) produced robust BOLD activation to synthetic whole face stimuli, but also to the internal facial features and head outlines. BOLD responses to whole face stimuli in FFA were significantly reduced after adaptation to whole faces, but not after adaptation to features or head outlines, whereas activation to head outlines was reduced after adaptation to both whole faces and head outlines. OFA showed no significant adaptation effects for matching adaptation and test conditions, but did exhibit cross-adaptation between whole faces and head outlines. The internal face features did not produce any significant adaptation within either FFA or OFA. Our results are consistent with a model in which independent populations of whole face-, feature-, and head outline-tuned neurons exist within face-sensitive regions of human occipito-temporal cortex, which in turn would support tasks such as viewpoint processing, emotion classification, and identity discrimination. ■

INTRODUCTION

A network of regions in human neocortex, including the fusiform gyrus (fusiform face area, or FFA; Grill-Spector, Knouf, & Kanwisher, 2004; Kanwisher, McDermott, & Chun, 1997) and inferior occipital gyrus (occipital face area, or OFA; Gauthier et al., 2000; Haxby, Hoffman, & Gobbini, 2000), is more strongly activated by faces than other classes of objects. These regions, and FFA in particular, provide robust responses to a wide variety of face stimuli, including gray-scale and color photographs, two-tone (Mooney) images, cartoon faces, cat faces, and various schematic and synthetic faces (see Kanwisher & Yovel, 2006, for a recent review). The importance of FFA and OFA in face processing is confirmed by case studies of patients with damage to occipito-temporal cortex who exhibit impaired face processing capabilities, a condition known as acquired prosopagnosia (Bouvier & Engel, 2006; Barton, Press, Keenan, & O'Connor, 2002; Sergent & Signoret, 1992).

Although we know *where* faces are processed in the brain, it is still not clear *how* faces are represented. Interestingly, FFA activation was equally robust to whole faces and to faces with no eyes (Tong, Nakayama, Moscovitch, Weinrib, & Kanwisher, 2000). The eyes presented in isolation produced less activation than the whole face, but the BOLD response was significantly greater to eyes than to other objects (e.g., houses). These findings raise the question: Do FFA neurons encode only whole face stimuli? In this kind of framework, which we will call the homogeneous model, the face-sensitive neurons are only partially activated by components of the face, such as the internal features or head outline, because these stimuli contain only a portion of the information found in the whole face. An alternative explanation, which we will call the heterogeneous model, is that facial elements activate independent neural populations within FFA; the fMRI response to the whole face stimulus reflects the combined activity of neurons tuned to specific facial elements.

Support for the heterogeneous model of face processing comes from a variety of electrophysiology and neuroimaging studies conducted in human and nonhuman primates. Early physiological investigations in monkey inferior temporal cortex (IT) and superior temporal sulcus (STS) described face-selective cells that responded to different features or subsets of features (Perrett, Hietanen, Oram, & Benson, 1992; Yamane, Kaji, & Kawano, 1988; Perrett, Rolls, & Caan, 1982). More recent work found that a sample of highly face-selective neurons in monkey STS also responded to round and oval objects such as clocks and apples (Tsao, Freiwald, Tootell, & Livingstone, 2006), which are perceptually similar to a head outline. Electrophysiological recordings from surface electrodes in human cortex revealed sites in the fusiform gyrus and inferior occipital gyrus that responded significantly to circular stimuli such as polar and radial gratings (Allison, Puce, Spencer, &

York University, Toronto, Ontario, Canada
*Current address: McMaster University, Hamilton, Ontario, Canada.

McCarthy, 1999). Concentric gratings have been shown to elicit BOLD activation within the functionally defined FFA that was approximately half of that produced by whole face stimuli (Wilkinson et al., 2000), indicating a sensitivity for circular structure in face-sensitive cortex. Visual images composed of small fragments of faces that contained sufficient information for both detection and individuation of faces elicited significant activation in the right hemisphere fusiform gyrus (Nestor, Vettel, & Tarr, 2008). Finally, face stimuli that were manipulated with binocular disparity to appear either behind or in front of a set of bars (Nakayama, Shimojo, & Silverman, 1989) modulated amodal completion of the face behaviorally, but elicited equivalent activation in the middle fusiform gyrus (Harris & Aguirre, 2008). It is therefore possible that human FFA functions as a site where information about individual facial elements are processed separately and integrated into whole face representations.

Although the nature of face representation in FFA is uncertain, even less is known about the role of OFA in face perception. It has been suggested that OFA is important for the representation of facial features prior to "holistic" face processing by FFA (Liu, Harris, & Kanwisher, 2002; Haxby et al., 2000). Unlike FFA, which is highly sensitive to the prototypical face configuration (i.e., high contrast elements placed within the top half of a curvilinear contour), OFA responses do not differentiate between square and circular outlines, or whether the high contrast elements were concentrated in the upper or lower half of the stimulus (Caldara et al., 2006). OFA also shows a greater bias toward the center of the visual field compared to FFA, suggesting that the cortical overrepresentation of the fovea may help in performing subtle discrimination tasks (Levy, Hasson, Avidan, Hendler, & Malach, 2001). Transcranial magnetic stimulation of right OFA was shown to disrupt the ability to identify specific facial features while the information about the relative spacing of those features was preserved (Pitcher, Walsh, Yovel, & Duchaine, 2007). However, other fMRI work suggests that OFA is also involved in whole face processing (Schiltz & Rossion, 2006), consistent with a single population of neurons that encode whole faces. Damage to OFA can severely disrupt face recognition, even when right hemisphere FFA is preserved (Steeves et al., 2006; Rossion et al., 2003). Furthermore, damage to OFA may interfere with the ability to use optimal information contained in the eye region to discriminate faces, suggesting that OFA may be critically involved in fine-grained distinctions between faces in identification tasks (Rossion, 2008; Caldara et al., 2005). It has also been proposed that OFA is involved in face detection tasks (Nestor et al., 2008). A better understanding of the nature of face representation in OFA is therefore critical to any model of face processing in the brain.

We used an event-related fMRI adaptation paradigm (Grill-Spector & Malach, 2001) to investigate whether these face-sensitive regions are solely tuned for whole faces, or whether they contain heterogeneous populations of neurons tuned to individual components of the face, as suggested by physiological investigations in primates (Tsao et al., 2006; Perrett et al., 1982). The neuronal response of a particular population to its preferred stimulus decreases during the adaptation phase, causing a decrease in the amplitude of that voxel's response during the test phase. If, however, two stimuli that activate independent neural populations are presented in the adaptation and test phases, there will be no adaptation, and subsequently, no reduction in the BOLD response (see Krekelberg, Boynton, & van Wezel, 2006, for a review). In this way, we can determine whether the face-sensitive regions of the brain are tuned to whole faces, as per the homogeneous model, or whether these regions of interest consist of separate populations of whole face, features, and outline-sensitive neurons, as per the heterogeneous model.

METHODS

Participants

Thirteen right-handed participants (6 women; mean age = 27.23 years, σ = 4.04) were recruited from the Center for Vision Research in Toronto and the Greater Hamilton Area. All participants had normal or corrected-to-normal vision and were free from any ocular or visual pathology.

fMRI Data Acquisition and Preprocessing

Data were acquired with the research -3T short bore GE Excite-HD magnet equipped with a customized eight-channel head coil at the Imaging Research Centre, St. Joseph's Hospital, Hamilton, Ontario. The experiment was conducted in accordance with the guidelines of the St. Joseph's Healthcare Research Ethics Board. Participants provided informed consent and were remunerated $25/hour for their time.

High-resolution ($0.5 \times 0.5 \times 0.8$ mm) 3-D anatomical images were acquired in the axial plane using a FastIR prep, SPGR whole-brain anatomical scan (Zip512, T1-weighted, flip angle = 12°, FOV = 24 cm, TE = 2.1 msec). Functional 2-D images were collected in the axial plane using a series of T2*-weighted gradient-echo (EPI) scans (TE = 35 msec, TR = 1250 msec, flip angle = 90°, FOV = 24 cm, interleaved acquisition, zero gap, in-plane resolution = 3.75×3.75 mm) to maximize BOLD contrast. Each scan typically contained 18 to 22 slices (4.0 mm thick) that extended from the top of the corpus callosum to the bottom of the temporal lobe.

A scanning session consisted of nine experimental functional scans, two localizer functional scans, and a high-resolution structural scan. Each scan lasted between 6 and 8 min, and the entire session was completed in 2 hr.

The MRI data were imported into BrainVoyager QX (v 1.9). Functional scans were preprocessed with BrainVoyager's

slice scan-time correction (cubic spline interpolation), linear trend removal, and motion correction algorithms. The FFA localizer scans were temporally smoothed with a 0.0093-Hz (3 cycles/scan) high-pass filter; event-related scans were not temporally smoothed. All functional scans were transformed into 4-D volumetric time courses (resampled to $3 \times 3 \times 3$ mm voxel resolution) and aligned to the high-resolution anatomical scans.

Stimulus Display

Visual stimuli were projected onto a screen attached to the head coil and viewed by a 45° angle mirror (MRIx Synchronization Control System; Thulborn & Associates, Bannockburn, IL). The viewing distance was approximately 30 cm, depending on the size of the participant's head. The 100% whole face and outline stimuli subtended on average $7.7° \times 10.3°$ of visual angle; features subtended approximately $5.1° \times 4.9°$.
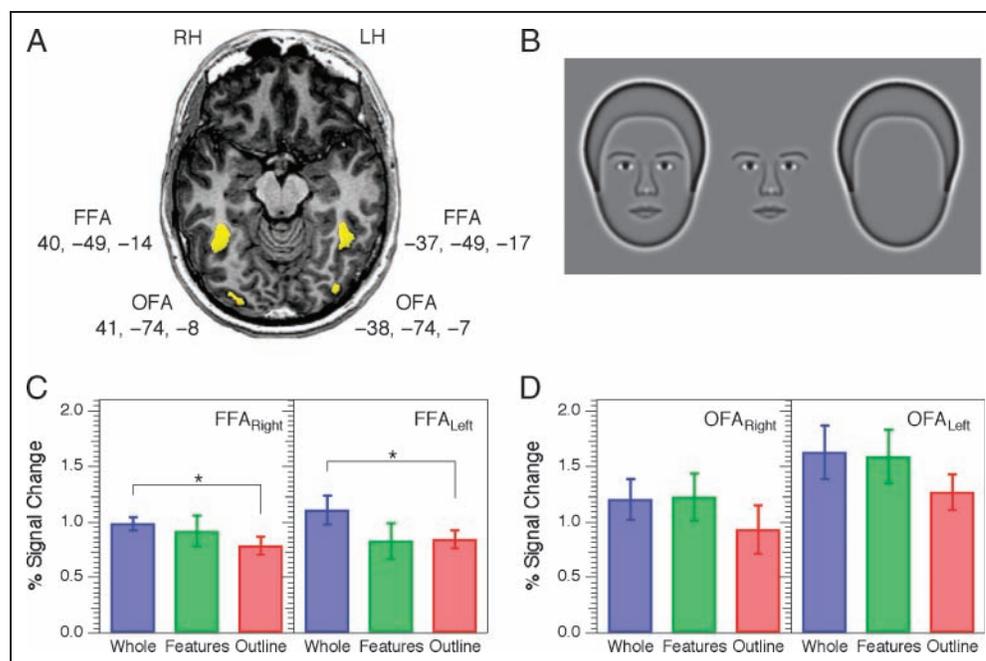
Stimuli

The synthetic face stimuli (Figure 1B) were generated using the procedures developed by Wilson, Loffler, and Wilkinson (2002). Briefly, the faces were derived from digital photographs of emotionally neutral, individual faces taken from the frontal view. Thirty-seven facial coordinates, measured relative to the bridge of the nose, were recorded for each face. The head and hairline coordinates were converted into sums of radial frequencies (Wilkinson, Wilson, & Habak, 1998), such that 23 num-

bers relative to the mean head radius described the head and hairline shape. A total of 14 polar coordinates described the locations of the eyes (x and y locations), eyebrows (position relative to eyes), nose (tip location, width), and mouth (left, right, and center locations; upper and lower lip thickness). An individual face was therefore represented as 37-dimensional vector. Although the features were generic, the positions of the eyes and eyebrows were individually specified, as were the locations, widths, and lengths of the nose and lips. All textural information of the faces was discarded. An average male face and an average female face were generated from the mean of the 37-dimensional vectors of two groups of seven faces. The individual faces were then normalized so that they varied from the group mean by 15%. The Gram–Schmidt orthogonalization procedure was used to ensure that all the faces within a group were orthogonal to each other in 37-dimensional face space.

In order to preserve local stimulus contrast, we constructed the features and outline stimuli separately. The stimuli were filtered with a radially symmetric difference of Gaussians with a peak frequency of 10.0 cycles per mean face width and a 2.0-octave bandwidth, in keeping with psychophysical estimates of the range of spatial frequencies over which optimal face identification occurs (Gold, Bennett, & Sekuler, 1999; Näsänen, 1999; Costen, Parker, & Craw, 1996; Hayes, Morrone, & Burr, 1986; Fiorentini, Maffei, & Sandini, 1983). We then simply added the features and outline images together to create the whole face stimuli. Stimulus contrast was held at 100% throughout the experiment. The images were sequentially compiled

Figure 1. (A) FFA in the middle fusiform gyrus and OFA in the inferior occipital gyrus were independently localized as regions that responded significantly greater to photographs of faces than to photographs of houses. Mean Talairach coordinates are indicated below the labels. (B) The synthetic face stimuli (Wilson et al., 2002) consisted of the whole face, features alone, or head outline alone. We used a set of eight male and eight female faces. (C) The activation in right and left FFA to the unadapted condition indicates robust responses to the whole face, features, and outline stimuli. (D) Similarly, bilateral OFA showed strong responses to the three different stimulus types. Error bars represent the standard error of the mean. Asterisks indicate significant differences for planned orthogonal comparisons ($p < .05$). Note that the data in C and D are the same as the unadapted responses shown in Figure 3.

using Matlab (v. 7) to create movie files that were then displayed through the MRIx stimulus display system at the scanner.

Face-sensitive Regions of Interest

Face-sensitive FFA and OFA (Figure 1A) were functionally defined using randomly ordered blocks of gray-scale face and house photographs previously employed by this laboratory (Loffler, Yourganov, Wilkinson, & Wilson, 2005). A general linear model analysis was used to determine which voxels sustained significantly greater responses to faces compared to houses (Kanwisher et al., 1997). Significantly greater activation for the face stimuli was evident in several areas of visual cortex, including the middle fusiform gyrus (FFA) and the inferior occipital gyrus (OFA). Fusiform activation was bilateral in all observers. Right inferior occipital activation was observed in 10 of 11 participants; left inferior occipital activation was found in 8 of 11 observers. The mean FFA and OFA Talairach coordinates are reported in Figure 1 for reference. However, it is important to note that the regions of interest were defined individually on nontransformed brains. Other regions, including the superior temporal sulcus, inferior temporal sulcus, and medial frontal gyrus, were identified in (nonoverlapping) subsets of participants, but did not produce meaningful or reliable time-course data in response to the adaptation paradigm, and were therefore not included in the analysis.

Retinotopic Mapping and Object-sensitive Regions of Interest

Eight participants returned for an additional day of scanning to obtain retinotopic maps. The functional and anatomical data from the additional scanning session were aligned to the main experimental dataset. Retinotopic areas V1, V2, V3, hV4, V3A/B, and LO1 were delineated according the BOLD phase maps elicited by standard rotating wedge and expanding ring stimuli (Larsson & Heeger, 2006; Engel, Glover, & Wandell, 1997).

Experimental Design

The main experimental scans employed a single-trial event-related design. A single event consisted of a 5-sec adaptation phase, a 5-sec test phase, and a 15-sec relaxation phase in which the BOLD response was allowed to return to baseline (Figure 2A). The participants were required to perform a two-interval forced-choice size discrimination task during both the adaptation and test phases to ensure (a) visual attention during both phases of the experiment, and (b) that differences in activation to the adaptation and test phases could not be attributed to differences in task requirements. Although it has been shown that the FFA response to faces is invariant up to a 200-fold increase in stimulus size (Andrews & Ewbank, 2004), we kept the size difference to a minimum (8%), which was large enough to produce good performance in the task (94% mean accuracy in all stimulus conditions), but also small enough to require strict visual attention.

Incorporating the task into both phases posed a challenge, however, in that we needed to find a way to have two presentations of the same stimulus (with a subtle size change) during the test phase without incurring further adaptation. In the test phase, we did not present the stimulus continuously for 5 sec. The first TR (1.25 sec) was a blank stimulus, which was a sufficiently long duration to avoid any masking by the adaptation stimulus, but still short enough to retain adaptation effects on the subsequent stimulus presentation. We then showed a single stimulus for one TR, a blank interstimulus interval (ISI) of one TR, and then the final stimulus presentation for one TR. The ISI of one TR is two to four times longer than the intervals commonly used in fast ER designs (typical ISI values range from 100 to 600 msec; e.g., Kourtzi & Huberle, 2005; Kourtzi, Tolias, Altmann, Augath, & Logothetis, 2003; Kourtzi & Kanwisher, 2000, 2001), therefore minimizing fast adaptation effects during the test phase. It has been suggested that short-interval adaptation fatigues the synaptic inputs to a cortical region, whereas long-interval adaptation may affect within-region processing (Epstein, Parker, & Feiler, 2008). Here we are interested in the effects of long-interval adaptation.

Four adaptation conditions (unadapted, whole face, features, and outline) and one of three test conditions (whole face, features, or outline) were presented within each adaptation scan. The face identity was held constant within a single event. Each adaptation scan contained 17 events and 16 identities. The event order was pseudorandomized such that each adaptation condition was preceded by all of the other adaptation conditions exactly once. The first event did not have a trial history, and was not analyzed; as such, each adaptation condition was presented four times within the scan. The first event always consisted of the mean female or male face, and was the single repeated identity within the scan. Participants completed three scans of each test condition, for a total of nine experimental scans.

Data Analysis

The mean percentage BOLD signal change from fixation baseline was calculated within the individually defined FFA and OFA regions of interest for each adaptation and test condition in all observers. The height of the peak time point associated with the unadapted condition was used as a measure of the BOLD signal for each adaptation condition. The peak time point for each unadapted response in each individual ROI was determined for every participant (range: 12.25 to 15 sec after the trial onset; $\sigma = 1.037$ sec). To assess whether the peak time points varied across testing conditions and ROIs, the time points

Figure 2. (A) The single-trial fMRI adaptation paradigm was incorporated in a two-alternative, forced-choice size discrimination task in both adaptation and test phases of the experiment. The participants were required to indicate which interval contained the 8% larger stimulus. The same facial identity was used within a single trial. During the relaxation phase, a fixation point appeared on the screen. BOLD time-course data from the right (B) and left (C) hemisphere FFA are shown for the three different adapting conditions. The gray vertical bar indicates the peak time point used to compare adaptation effects.
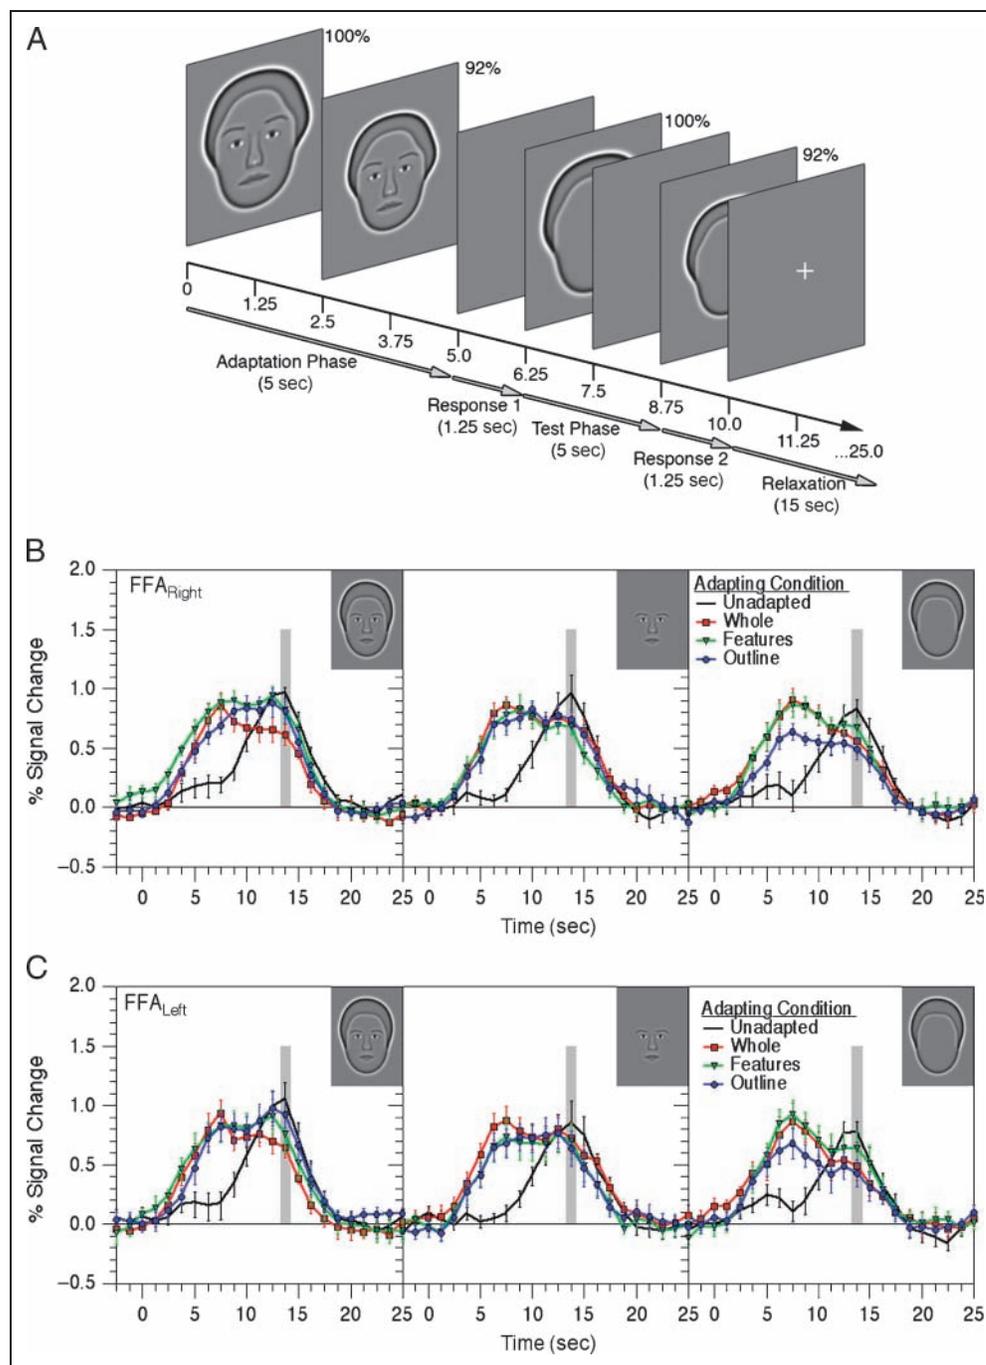
for those participants with all four ROIs were submitted to an ROI × Hemisphere × Test Condition repeated measures ANOVA. Interestingly, the peaks were significantly earlier in FFA than OFA, as indicated by a main effect of ROI [FFA peak = 13.125; OFA peak = 13.359; mean difference = -0.26 sec, $SE = 0.094$; $F(1, 7) = 7.609$, $p < .05$]. However, there was no main effect of hemisphere or test condition, and no significant interaction between any of the variables. Therefore, within a given ROI, we defined the peak time point for each participant as the mean of the peaks across the three testing conditions, rounded to the nearest whole number.

Once the peak time point was determined, the percent signal change for each adaptation and test condition, averaged across all voxels within the ROI, was recorded. Repeated measures ANOVAs were conducted on the signals for each test condition. In the case where an ANOVA violated the assumption of sphericity, we implemented the Greenhouse–Geisser epsilon (ε) correction to control for Type I error. Bonferroni-corrected multiple post hoc comparisons were used to assess significance of the different adapting stimuli. We were mainly interested in the comparison of the different adapting conditions to the unadapted response, but still used the Bonferroni correction

based on the total number of possible multiple comparisons. In one participant, the signal from left OFA was not statistically significant from zero in the unadapted response to whole faces and was therefore not included in the left OFA group analysis.

An analysis using the three adjacent peak time points of the unadapted condition produced the same results for left FFA, right FFA, and right OFA as the analysis with the single peak time point.

RESULTS

Unadapted Response to Whole Faces, Features, and Outlines in FFA and OFA

Synthetic whole faces, facial features, and head outlines in the unadapted test condition (Figure 1B) all elicited robust BOLD activation from FFA and OFA. Activation in FFA was significantly greater for the whole faces than for head outlines in both right [$F(1, 10) = 6.72, p < .05$] and left [$F(1, 10) = 6.96, p < .05$] hemispheres, with intermediate activation to facial features (Figure 1C). Both right and left OFA regions responded equally well to whole faces, features, and outlines (Figure 1D), suggesting that the different parts of a face were equally salient to these face-sensitive neurons. Although responses in OFA were more variable than in FFA, OFA activation was of the same magnitude as that observed in FFA, confirming a crucial role for OFA in face representation (Steeves et al., 2006; Rossion et al., 2003; Haxby et al., 2000). Alone, these results do not distinguish between the heterogeneous and homogenous models of face representation in either region. The fMRI adaptation paradigm is a powerful tool to distinguish between these two hypotheses, and may shed some light on the separate functions that FFA and OFA play in face processing.

fMRI Adaptation within FFA

We used a single-trial event-related fMRI adaptation paradigm (Figure 2A; Grill-Spector & Malach, 2001) to distinguish between the two equally plausible models of face representation in visual cortex. Within an individual 7-min scan, participants adapted to whole faces, features, outlines, or a blank screen of mean luminance (the unadapted condition) while the test stimulus was held constant. Mean BOLD time courses from right and left FFA are shown in Figures 2B and 2C. To assess the effects of adaptation in the FFA region of interest, we measured the difference between the peak of the unadapted BOLD time course and the corresponding point of the adapted time course during the test phase.

In the case where test stimuli consisted of whole faces, a significant main effect of adapting condition was found in right [$F(3, 30) = 12.567, p < .01$] and left [$F(3, 30) = 8.822, p < .01$] FFA (Figure 3A, B). Consistent with previous fMRI adaptation experiments (Fang, Murray, & He,

2007; Mazard, Schiltz, & Rossion, 2006; Loffler et al., 2005; Yovel & Kanwisher, 2005; Andrews & Ewbank, 2004; Winston, Henson, Fine-Goulden, & Dolan, 2004; Grill-Spector & Malach, 2001; Gauthier et al., 2000), the presentation of whole face stimuli in the adaptation phase significantly reduced the response to the whole face relative to the unadapted response in both right and left FFA [unadapted–whole: right FFA, $t(10) = 6.783, p < .01$; left FFA, $t(10) = 6.440, p < .01$]. However, activation to whole face stimuli was not significantly reduced after adaptation to either features or head outlines. An effect of adapting condition in the features test condition was marginally significant in right FFA only [$F(3, 30) = 3.502, p = .060, \varepsilon = 0.560$]. Surprisingly, none of the post hoc comparisons were significant; adaptation to any of the three stimulus types did not significantly reduce the response to features, even when the features were the adapting stimuli. The outline test condition showed significant overall effects of adaptation stimulus [right FFA, $F(3, 30) = 13.875, p < .01$; left FFA, $F(3, 30) = 6.973, p < .01$]. Pairwise comparisons showed that adaptation to both whole faces and head outlines substantially reduced the response to head outlines [unadapted–whole: right FFA, $t(10) = 4.097, p < .05$; left FFA, $t(10) = 3.549, p < .05$; unadapted–outline: right FFA, $t(10) = 5.596, p < .01$; left FFA, $t(10) = 3.670, p < .05$]. In sum, whole faces and outlines were self-adapting, and adaptation to whole faces also caused reduced activation to outlines.

fMRI Adaptation within OFA

An identical analysis yielded a different pattern of responses to whole faces, features, and outlines in the OFA region of interest (Figure 3C, D). Most notably, fewer significant adaptation effects were observed in OFA compared to FFA. In both right and left OFA, the overall effect of adapting stimulus reached significance for the whole face test condition [right OFA, $F(3, 27) = 4.536, p = .05$; left OFA, $F(3, 18) = 3.828, p < .05$]. Adaptation to the head outlines significantly reduced the response to whole faces in left OFA [unadapted–outline: $t(6) = 4.506, p < .05$]. A significant effect of adaptation was seen in the features condition in left OFA [$F(3, 18) = 3.789, p < .05$], however, none of the multiple comparisons were robust. The strongest adaptation effects were seen in the head outline test condition, where both right and left OFA yielded significant main effects of adapting condition [right OFA, $F(3, 27) = 5.085, p < .05, \varepsilon = 0.604$; left OFA, $F(3, 18) = 5.976, p < .01$]. Adaptation to whole faces significantly reduced the responses to head outlines in both left [unadapted–whole: $t(6) = 3.910, p < .05$] and right OFA [unadapted–whole: $t(9) = 3.369, p < .05$]. An additional adaptation effect was observed in left OFA (Figure 3D), in which the head outlines significantly adapted to themselves [unadapted–outline: $t(6) = 5.884, p < .01$]. The somewhat different pattern of adaptation, particularly for the whole face and outline conditions, suggests that OFA

Figure 3. Peak time point data from (A) right FFA, (B) left FFA, (C) right OFA, and (D) left OFA are shown for the four adapting conditions (unadapted, whole, features, and outline) and three different test conditions (whole, features, and outline). Error bars indicate standard error of the mean. Asterisks indicate significant Bonferroni-corrected pairwise comparisons at the .05 (*) and .01 (**) alpha levels.
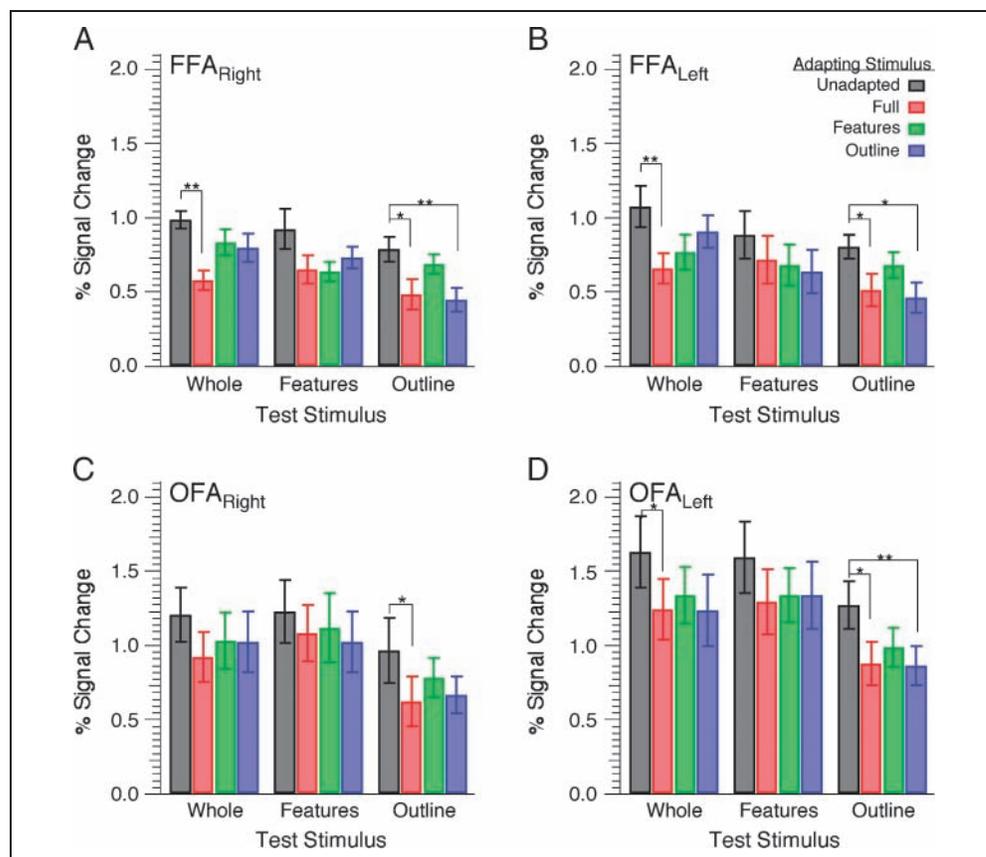
plays a different role in the face processing network than the fusiform face-sensitive neurons, and also indicates potential hemispheric differences.

Early Visual Areas

Although our stimuli varied in size by 8% within individual trials (Figure 2A), adaptation in the early visual areas could potentially propagate throughout the visual processing hierarchy. It is therefore extremely important to ensure that our observed adaptation effects were not simply inherited from regions of visual cortex that were not isolated by our FFA/OFA localizer scans. Following the definitions provided by Larsson and Heeger (2006), we mapped the BOLD phases produced by rotating wedge and expanding ring stimuli to delineate areas V1, V2, V3, V3A/B, hV4, and LO1. Unlike FFA and OFA, no significant adaptation was observed in any of the adapting or test conditions (Figure 4), indicating that low-level adaptation effects were not responsible for the patterns of activation observed in the face-sensitive occipito-temporal regions of interest.

DISCUSSION

The fMRI adaptation paradigm was designed to specifically test two plausible models of face encoding in human occipito-temporal cortex, the homogeneous and heterogeneous models described above. For several reasons, our data are *not* consistent with a homogeneous population of face-tuned neurons in FFA. First, even though all three types of stimuli elicited robust activation from FFA, there was relatively little cross-adaptation between the different stimuli. If a single population of face-tuned neurons was partially activated by the features and head outlines, it would be expected that adaptation to features and outlines would reduce the BOLD response to subsequent presentations of the whole face. The results shown in Figures 3A and 3B indicate that the only effective adapting stimulus in the whole face test condition was the whole face stimulus. Second, one would predict that adaptation to the internal features would fatigue the face-tuned neurons, which would in turn show a reduced response to the presentation of the head outline. We found no cross adaptation between features and head outlines, regardless of the order of adaptation and testing, consistent with independent representations of these two stimulus types. Third, the asymmetrical pattern of adaptation to whole faces and outlines, that is, adaptation to whole faces reduced the response to outlines even though adaptation to outlines did not reduce the response to whole faces, suggests that different neural populations are activated by the whole face and head outline stimuli. For these reasons, our results are more consistent with the heterogeneous

model in which FFA is composed of separate populations of cells tuned to whole faces, internal facial features, and head shape.

A schematic of the possible organization of FFA neurons that accounts for the observed pattern of fMRI adaptation is depicted in Figure 5A. In this framework, neurons that are tuned to whole faces do not receive direct input from other visual areas. Rather, face information is received by neurons tuned to internal facial features and head outlines, which then project to whole face cells. Importantly, the whole face cells require inputs from both feature and outline cells to fire (and therefore adapt), the equivalent of a logical AND operation or neural threshold. The existence of such super-additive face-selective AND neurons has been reported in macaque anterior inferotemporal cortex (Kobatake & Tanaka, 1994). Similarly, human fMRI studies have demonstrated that right FFA responds maximally to schematic face stimuli that contain high contrast internal elements in the top half of the stimulus placed within a curvilinear contour. Right hemisphere FFA activation, and also the behavioral "faceness" ratings of the stimulus, dropped significantly if the elements were (a) presented within a square outline, (b) shifted to the lower half of the stimulus, or (c) asymmetrically arranged across the vertical axis of the stimulus (Caldara & Seghier, 2009; Caldara et al., 2006). These findings, as well as the observed pattern of adaptation in the present study, provide converging evidence of a population of neurons in the fusiform gyrus that respond to the conjunction of internal facial features and head outlines.

To demonstrate its plausibility, a very simple quantitative version of this FFA model was implemented. The model comprises feature neurons (FE), outline neurons (OU), and whole face (WF) neurons. FE neurons gave response 1.0 whenever features were present in the stimulus, and OU neurons likewise gave an unadapted response of 1.0 in response to a head outline. To implement the AND operation, a threshold $\theta = 1.0$ was subtracted from the summed inputs to WF. Thus, WF = 0 when only FE or OU is active, but WF = 1 when both are activated by a full face. Following adaptation, both OU and WF were assumed to generate reduced outputs scaled by the adaptation factor A = 0.7. As there was no statistically significant fMRI adaptation shown for features, FE was assumed not to adapt at all. As shown in Table 1, the interesting model adaptation results concern adaptation of WF and OU neurons. In agreement with the data, neither adapts significantly to features, but both are adapted significantly by whole faces. In addition, OU are adapted by outlines, which also agrees with the data. The one discrepancy is the model prediction of WF adaptation by outlines (italicized in Table 1), an effect that was absent in the fMRI responses (Figure 3A, B). Closer inspection of the fMRI data revealed that the WF response did show a trend toward adaptation by outlines ($p = .125$). Thus, the model does correctly predict that WF responses should be greatest after feature adaptation, smaller after outline adaptation, and smallest after whole face adaptation. The explanation for this pattern is that both the OU and WF neurons adapt when stimulated, and whole faces adapt both, whereas
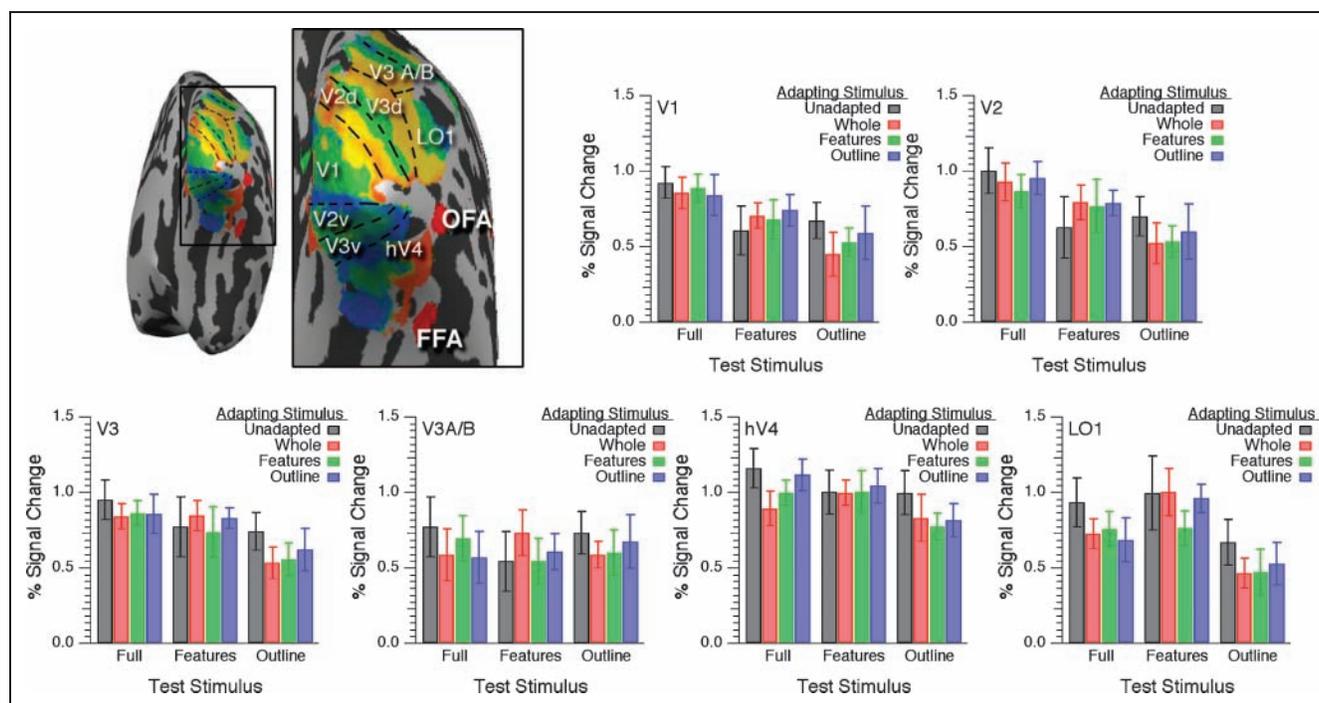

Figure 4. The percent signal change of the peak time points are shown for the retinotopically defined regions of visual cortex in eight participants. There were no significant adaptation effects in any of these early visual areas.
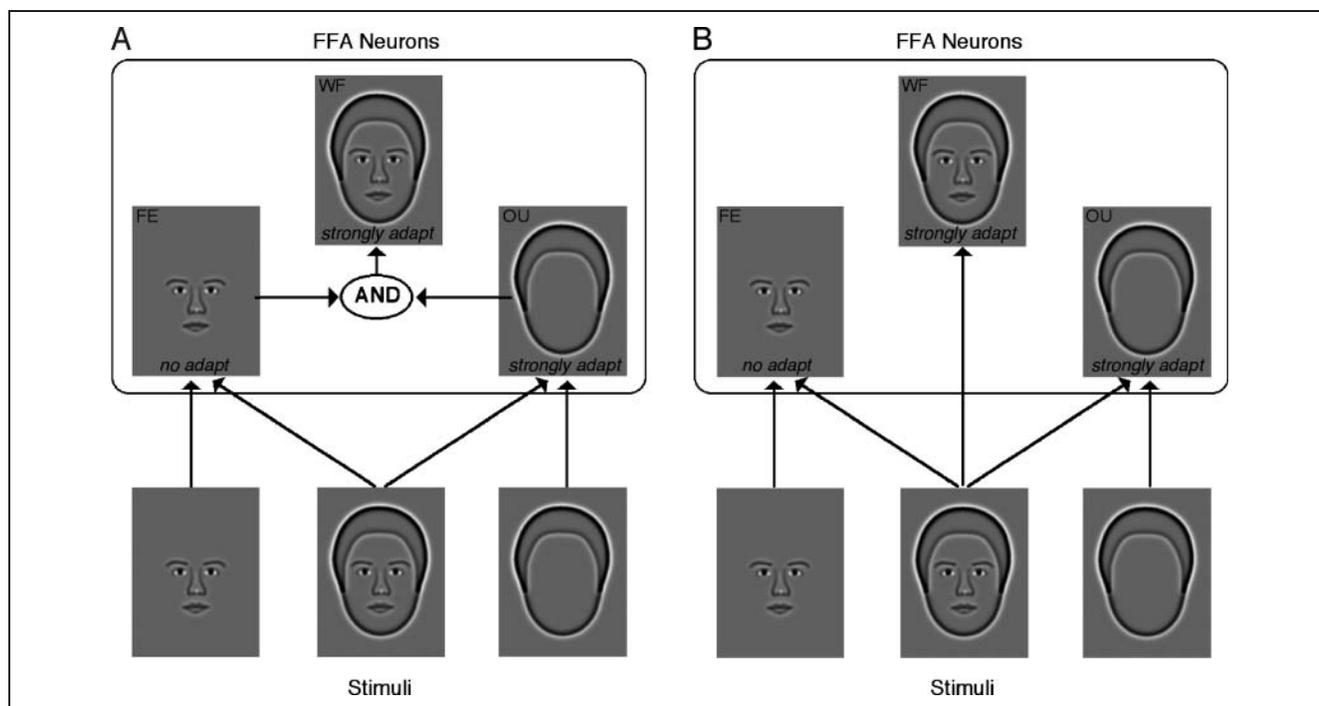
Figure 5. Two possible schematic representations of whole face-, features-, and outline-tuned neurons in FFA. (A) Neurons tuned to whole faces (WF) require inputs from both features and outline cells to fire. Whole face and outline (OU) neurons demonstrate adaptation, whereas features neurons (FE) produce consistent firing even after repeated presentations of identical stimuli. (B) An independent pathway activates populations of neurons selective for whole face stimuli. OU neurons are adapted by both whole face and outline stimuli, whereas FE neurons do not adapt to either features or whole faces.

outlines adapt only OU. In sum, this simple model accurately replicates the unadapted responses and the adapted responses under all conditions except one, and the data do show a trend toward adaptation in that case. A more complex model would likely rectify this one discrepancy.

As an alternative to the model in Figure 5A, an anonymous reviewer suggested the model depicted in Figure 5B. On the assumption that OU and WF neurons adapt, whereas FE neurons do not, this model can explain all of our statistically significant adaptation data. Note, however, that this model does not incorporate any computations among the three neural populations in FFA, which seems less plausible than the previous model that includes the AND operation. Indeed, the first model is consistent with a sequential construction of face representations with a final stage in FFA. The latter model implies that face construction is already complete at a level prior to FFA, which leaves it unclear why features and head outlines should be independently represented in FFA. Further research will be required to untangle these issues.

Why did the features stimuli fail to adapt, even to themselves? The robust size constancy that has been consistently demonstrated in FFA (Andrews & Ewbank, 2004; Grill-Spector & Malach, 2001; Kanwisher et al., 1997) suggests that the subtle stimulus size variation implemented for the behavioral task was insufficient to disrupt adaptation, at least for whole face and outline stimuli. However, it may be advantageous for the feature-encoding mecha-

nisms within FFA to maintain a high degree of excitability, which could be implemented as either a strong resistance to adaptation or a quick recovery from adaptation. It has been suggested that FFA is biased for processing foveal regions of the visual field, which would presumably increase sensitivity to fine details in the stimulus (Levy et al., 2001). Indeed, FFA is extremely sensitive to subtle changes in internal feature spacing (Nestor et al., 2008; Yovel & Kanwisher, 2004). Given that internal facial features, particularly the eyebrows and mouth, are highly mobile in naturalistic situations, adaptation to specific facial features, or a slow recovery from adaptation, may actually compromise FFA's face discrimination abilities. The sluggish nature of the BOLD response does not lend itself to measuring quickly fluctuating neural responses; adaptation effects to features may be uncovered using other neuroimaging methods with greater temporal resolution, such as MEG

Table 1. Quantitative Outputs from the Heterogeneous Model of Face Processing in Figure 5A

	Adapt Whole Face	Adapt Features	Adapt Outline
WF response	0.49	1.0	*0.7*
OU response	0.7	1.0	0.7
FE response	1.0	1.0	1.0

(e.g., Liu et al., 2002), a combined EEG/fMRI approach, or using single-cell physiology in an animal model.

It is also conceivable that adaptation to internal facial features could occur in other regions of the face processing network. However, our results suggest that the adaptation does not occur early on in the visual processing hierarchy in areas such as OFA. It is generally held that FFA is crucial for face identification and discrimination, whereas activity in STS is linked with processing changeable aspects of faces, including facial expression, viewpoint, and gaze (Andrews & Ewbank, 2004; Winston et al., 2004; Haxby et al., 2000). Although we were unable to reliably localize face-sensitive areas beyond FFA and OFA, a recent study suggests that the use of dynamic, rather than static, images of faces and objects may elicit more consistent activation from a wide range of brain regions implicated in face perception (Fox, Iaria, & Barton, 2008). Whether more robust adaptation to internal facial features, or cross adaptation between whole faces, features, and head outlines, occurs in a different region of interest in the face processing network, such as STS, middle temporal gyrus, middle frontal gyrus, posterior cingulate gyrus, or amygdala (Maurer et al., 2007; Andrews & Ewbank, 2004; Winston et al., 2004; Morris, DeGelder, Weiskrantz, & Dolan, 2001; Haxby et al., 2000), remains an open question.

Stimulus face familiarity may have also played a role in determining the extent of adaptation to the internal facial features. Behavioral studies have shown that observers are more able to match isolated facial features to the whole face image for familiar than novel faces (Clutterbuck & Johnston, 2002; Young, Hay, McWeeny, Flude, & Ellis, 1985; Ellis, Shepherd, & Davies, 1979). In general, face familiarity seems to increase the invariance of the activation in FFA to different views of whole faces (Ewbank & Andrews, 2008; Pourtois, Schwartz, Seghier, Lazeyras, & Vuilleumier, 2005). It is therefore conceivable that the representation of the internal features may become more robust, and therefore, exhibit greater adaptation effects with increased familiarity. However, it has also been suggested that the effect of familiarity may interact with the holistic face processing. In a recent experiment, Harris and Aguirre (2008) used facial stimuli that were partially occluded by a series of horizontal bars. The bars were manipulated in stereoscopic vision to appear either in front of the faces (in which case, the faces were completed and perceived as wholes) or behind the faces (in which case, the faces were not processed holistically). Adaptation in the right fusiform gyrus was greater for familiar than unfamiliar faces when the stimuli were holistically processed, but when the facial features were not integrated into holistic representations, unfamiliar faces elicited greater adaptation (Harris & Aguirre, 2008). Interestingly, the interaction between familiarity and holistic processing was not observed in the left fusiform gyrus, although there was a slight trend toward reduced adaptation for holistically processed faces. In the context of the present experiment, familiarity with the facial stimuli may increase the strength of adaptation to whole faces, but it is not guaranteed that the same effect would automatically generalize to the internal facial features.

The results from the OFA analysis produced markedly different activation patterns than those observed in FFA. OFA showed strong activation to all three stimulus types, but fewer conditions elicited significant adaptation compared to FFA. Our findings are similar to previous reports that the extent of adaptation to whole face stimuli is less in OFA than in FFA (Mazard et al., 2006; Yovel & Kanwisher, 2005; Andrews & Ewbank, 2004). The OFA response is also more sensitive to the retinal position of the face (Kovacs, Cziraki, Vidnyanszky, Schweinberger, & Greenlee, 2008), as well as to subtle changes in the spacing of internal features (Rotshtein, Vuilleumier, Winston, Driver, & Dolan, 2007; Yovel & Kanwisher, 2005). However, the significant cross-adaptation effects between whole faces and outlines, as well as the self-adaptation to outlines in left OFA, suggest that OFA is sensitive to global head shape. Perhaps these findings are not a total surprise, given the close proximity of OFA to brain regions associated with object and contour processing (the lateral occipital complex; Grill-Spector, Kourtzi, & Kanwisher, 2001).

The pattern of adaptation in the OFA is difficult to reconcile with a single homogenous population of face-tuned neurons. As with FFA, our results are more consistent with a heterogeneous neuronal population within OFA that collectively produces activation to a range of facial stimuli. It has recently been proposed that strong reciprocal connectivity between OFA and FFA acts to refine the representation of the face image (Rossion, 2008); although consistent with case studies of prosopagnosia (Schiltz & Rossion, 2006; Steeves et al., 2006), further research is required to determine if such a model operates within the normal adult population.

It must be noted that a lack of adaptation to a stimulus in a particular region of interest does not necessarily indicate a lack of neurons tuned to that particular stimulus. For example, up until very recently, fMRI adaptation techniques could not detect orientation tuning in V1 without prolonged adaptation periods of 20 sec or more (e.g., Boynton & Finney, 2003), despite the fact that orientation selectivity is a well-known property of V1 neurons. Adaptation to a particular orientation was sufficient to introduce a subtle bias across the V1 voxels, but the bias was too subtle to be detected using conventional statistical comparisons (i.e., subtraction). The implementation of more sophisticated analysis techniques (i.e., multi-voxel pattern classification and machine learning algorithms) allowed for these subtle biases to be revealed (Kamitani & Tong, 2005). In the context of the present experiment, it may be the case that the adaptation response to internal features is revealed through a more distributed pattern across the FFA and OFA regions that is not easily detected through the analysis employed here. In fact, our lab has recently utilized such pattern classification techniques (linear pattern classifiers obtained from

Support Vector Machines) to show that distinct spatial patterns of activation to whole faces, features, and outlines can be accurately discriminated in both FFA and OFA, using the data from the present experiment (Betts, Nichols, & Wilson, 2009). Additional experiments using a block design in an independent group of participants verified that the classification performance was robust. The pattern classification findings further support our interpretation of the present fMRI adaptation data, namely, that distinct distributed populations of neurons encode whole faces and their parts, and suggest different roles for FFA and OFA in the processing of whole faces and head outlines.

Several psychophysical adaptation studies have produced results consistent with dissociable neural populations that selectively encode various aspects of facial stimuli. For example, adaptation to upright faces produced strong gender and face distortion aftereffects in inverted faces, whereas adaptation to inverted faces produced relatively weak aftereffects in upright faces (Watson & Clifford, 2006; Rhodes et al., 2004). Watson and Clifford (2006) suggest that three separate neural populations that encode holistic information for upright faces, parts-based/featural information for upright faces, and featural information for inverted faces could account for their experimental results. Similarly, adaptation to eye position, identity strength, and masculinity produced significant aftereffects for faces of the same sex but did not transfer to faces of the opposite sex, implying separate neural populations for male and female faces (Little, DeBruine, & Jones, 2005). Little, DeBruine, Jones, and Waitt (2008) have shown similar category contingent aftereffects for age, race, and species, and propose that different categories of face stimuli are represented by distinct neural substrates in the visual system.

In summary, our fMRI adaptation data are consistent with the hypothesis that facial stimuli are encoded by independent populations of neurons in human occipito-temporal cortex tuned to whole faces, internal facial features, and global head shape. Furthermore, we suggest that the integration of facial features and head outlines into whole face representations occurs in FFA. This conceptual framework is consistent with physiological investigations in monkey STS that have revealed highly specific responses to whole faces as well as circles (Tsao et al., 2006), as well as neuroimaging work that showed strong responses to concentric patterns in human FFA (Wilkinson et al., 2000; Allison et al., 1999). In keeping with the present results, we hypothesize that these circular stimuli activated neurons tuned to head shape. Under our proposed model, objects that contain structural similarities to either head outlines (e.g., "Greebles" and "smoothies": Op de Beeck, Baker, DiCarlo, & Kanwisher, 2006; Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999) or internal facial features (e.g., cars: Grill-Spector et al., 2004) should be sufficient to activate FFA, but the best overall response will be provided by the presentation of both features *and* outlines. In short, our model provides an explanation for a phenomenon

that has intrigued the neuroimaging community for over a decade: Nothing activates the FFA like a whole face.

REFERENCES

Allison, T., Puce, A., Spencer, D. D., & McCarthy, G. (1999). Electrophysiological studies of human face perception: I. Potentials generated in occipitotemporal cortex by face and non-face stimuli. *Cerebral Cortex, 9,* 415–430.

Andrews, T. J., & Ewbank, M. P. (2004). Distinct representations for facial identity and changeable aspects of faces in the human temporal lobe. *Neuroimage, 23,* 905–913.

Barton, J. J., Press, D. Z., Keenan, J. P., & O'Connor, M. (2002). Lesions of the fusiform face area impair perception of facial configuration in prosopagnosia. *Neurology, 58,* 71–78.

Betts, L. R., Nichols, D. F., & Wilson, H. R. (2009). Classification of fMRI activation patterns in face-sensitive cortex to the parts and location of faces. [Abstract]. *Journal of Vision, 9,* 553a. doi:10.1167/9.8.553.

Bouvier, S. E., & Engel, S. A. (2006). Behavioral deficits and cortical damage loci in cerebral achromatopsia. *Cerebral Cortex, 16,* 183–191.

Boynton, G. M., & Finney, E. M. (2003). Orientation-specific adaptation in human visual cortex. *Journal of Neuroscience, 23,* 8781–8787.

Caldara, R., Schyns, P., Mayer, E., Smith, M. L., Gosselin, F., & Rossion, B. (2005). Does prosopagnosia take the eyes out of face representations? Evidence for a defect in representing diagnostic facial information following brain damage. *Journal of Cognitive Neuroscience, 17,* 1652–1666.

Caldara, R., & Seghier, M. L. (2009). The fusiform face area responds automatically to statistical regularities optimal for face categorization. *Human Brain Mapping, 30,* 1615–1625.

Caldara, R., Seghier, M. L., Rossion, B., Lazeyras, F., Michel, C., & Hauert, C. A. (2006). The fusiform face area is tuned for curvilinear patterns with more high-contrasted elements in the upper part. *Neuroimage, 31,* 313–319.

Clutterbuck, R., & Johnston, R. A. (2002). Exploring levels of face familiarity by using an indirect face-matching measure. *Perception, 31,* 985–994.

Costen, N. P., Parker, D. M., & Craw, I. (1996). Effects of high-pass and low-pass spatial filtering on face identification. *Perception & Psychophysics, 58,* 602–612.

Ellis, H. D., Shepherd, J. W., & Davies, G. M. (1979). Identification of familiar and unfamiliar faces from internal and external features: Some implications for theories of face recognition. *Perception, 8,* 431–439.

Engel, S. A., Glover, G. H., & Wandell, B. A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cerebral Cortex, 7,* 181–192.

Epstein, R. A., Parker, W. E., & Feiler, A. M. (2008). Two kinds of FMRI repetition suppression? Evidence for dissociable neural mechanisms. *Journal of Neurophysiology, 99,* 2877–2886.

Ewbank, M. P., & Andrews, T. J. (2008). Differential sensitivity for viewpoint between familiar and unfamiliar faces in human visual cortex. *Neuroimage, 40,* 1857–1870.

Fang, F., Murray, S. O., & He, S. (2007). Duration-dependent FMRI adaptation and distributed viewer-centered face representation in human visual cortex. *Cerebral Cortex, 17,* 1402–1411.

Fiorentini, A., Maffei, L., & Sandini, G. (1983). The role of high spatial frequencies in face perception. *Perception, 12,* 195–201.

Fox, C. J., Iaria, G., & Barton, J. J. (2008). Defining the face processing network: Optimization of the functional localizer in fMRI. *Human Brain Mapping, 30,* 1637–1651.

Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform "face area" increases with expertise in recognizing novel objects. *Nature Neuroscience, 2,* 568–573.

Gauthier, I., Tarr, M. J., Moylan, J., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). The fusiform "face area" is part of a network that processes faces at the individual level. *Journal of Cognitive Neuroscience, 12,* 495–504.

Gold, J., Bennett, P. J., & Sekuler, A. B. (1999). Identification of band-pass filtered letters and faces by human and ideal observers. *Vision Research, 39,* 3537–3560.

Grill-Spector, K., Knouf, N., & Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nature Neuroscience, 7,* 555–562.

Grill-Spector, K., Kourtzi, Z., & Kanwisher, N. (2001). The lateral occipital complex and its role in object recognition. *Vision Research, 41,* 1409–1422.

Grill-Spector, K., & Malach, R. (2001). fMR-adaptation: A tool for studying the functional properties of human cortical neurons. *Acta Psychologica (Amsterdam), 107,* 293–321.

Harris, A., & Aguirre, G. K. (2008). The representation of parts and wholes in face-selective cortex. *Journal of Cognitive Neuroscience, 20,* 863–878.

Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences, 4,* 223–233.

Hayes, T., Morrone, M. C., & Burr, D. C. (1986). Recognition of positive and negative bandpass-filtered images. *Perception, 15,* 595–602.

Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience, 8,* 679–685.

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,* 4302–4311.

Kanwisher, N., & Yovel, G. (2006). The fusiform face area: A cortical region specialized for the perception of faces. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences, 361,* 2109–2128.

Kobatake, E., & Tanaka, K. (1994). Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex. *Journal of Neurophysiology, 71,* 856–867.

Kourtzi, Z., & Huberle, E. (2005). Spatiotemporal characteristics of form analysis in the human visual cortex revealed by rapid event-related fMRI adaptation. *Neuroimage, 28,* 440–452.

Kourtzi, Z., & Kanwisher, N. (2000). Cortical regions involved in perceiving object shape. *Journal of Neuroscience, 20,* 3310–3318.

Kourtzi, Z., & Kanwisher, N. (2001). Representation of perceived object shape by the human lateral occipital complex. *Science, 293,* 1506–1509.

Kourtzi, Z., Tolias, A. S., Altmann, C. F., Augath, M., & Logothetis, N. K. (2003). Integration of local features into global shapes: Monkey and human fMRI studies. *Neuron, 37,* 333–346.

Kovacs, G., Cziraki, C., Vidnyanszky, Z., Schweinberger, S. R., & Greenlee, M. W. (2008). Position-specific and position-invariant face aftereffects reflect the adaptation of different cortical areas. *Neuroimage, 43,* 156–164.

Krekelberg, B., Boynton, G. M., & van Wezel, R. J. (2006). Adaptation: From single cells to BOLD signals. *Trends in Neurosciences, 29,* 250–256.

Larsson, J., & Heeger, D. J. (2006). Two retinotopic visual areas in human lateral occipital cortex. *Journal of Neuroscience, 26,* 13128–13142.

Levy, I., Hasson, U., Avidan, G., Hendler, T., & Malach, R. (2001). Center–periphery organization of human object areas. *Nature Neuroscience, 4,* 533–539.

Little, A. C., DeBruine, L. M., & Jones, B. C. (2005). Sex-contingent face after-effects suggest distinct neural populations code male and female faces. *Proceedings of the Royal Society of London, Series B, Biological Sciences, 272,* 2283–2287.

Little, A. C., DeBruine, L. M., Jones, B. C., & Waitt, C. (2008). Category contingent aftereffects for faces of different races, ages and species. *Cognition, 106,* 1537–1547.

Liu, J., Harris, A., & Kanwisher, N. (2002). Stages of processing in face perception: An MEG study. *Nature Neuroscience, 5,* 910–916.

Loffler, G., Yourganov, G., Wilkinson, F., & Wilson, H. R. (2005). fMRI evidence for the neural representation of faces. *Nature Neuroscience, 8,* 1386–1390.

Maurer, D., O'Craven, K. M., Le Grand, R., Mondloch, C. J., Springer, M. V., Lewis, T. L., et al. (2007). Neural correlates of processing facial identity based on features versus their spacing. *Neuropsychologia, 45,* 1438–1451.

Mazard, A., Schiltz, C., & Rossion, B. (2006). Recovery from adaptation to facial identity is larger for upright than inverted faces in the human occipito-temporal cortex. *Neuropsychologia, 44,* 912–922.

Morris, J. S., DeGelder, B., Weiskrantz, L., & Dolan, R. J. (2001). Differential extrageniculostriate and amygdala responses to presentation of emotional faces in a cortically blind field. *Brain, 124,* 1241–1252.

Nakayama, K., Shimojo, S., & Silverman, G. H. (1989). Stereoscopic depth: Its relation to image segmentation, grouping, and the recognition of occluded objects. *Perception, 18,* 55–68.

Näsänen, R. (1999). Spatial frequency bandwidth used in the recognition of facial images. *Vision Research, 39,* 3824–3833.

Nestor, A., Vettel, J. M., & Tarr, M. J. (2008). Task-specific codes for face recognition: How they shape the neural representation of features for detection and individuation. *PLoS ONE, 3,* e3978.

Op de Beeck, H. P., Baker, C. I., DiCarlo, J. J., & Kanwisher, N. G. (2006). Discrimination training alters object representations in human extrastriate cortex. *Journal of Neuroscience, 26,* 13025–13036.

Perrett, D. I., Hietanen, J. K., Oram, M. W., & Benson, P. J. (1992). Organization and functions of cells responsive to faces in the temporal cortex. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences, 335,* 23–30.

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.

Pitcher, D., Walsh, V., Yovel, G., & Duchaine, B. (2007). TMS evidence for the involvement of the right occipital face area in early face processing. *Current Biology, 17,* 1568–1573.

Pourtois, G., Schwartz, S., Seghier, M. L., Lazeyras, F., & Vuilleumier, P. (2005). View-independent coding of face identity in frontal and temporal cortices is modulated by familiarity: An event-related fMRI study. *Neuroimage, 24,* 1214–1224.

Rhodes, G., Jeffery, L., Watson, T. L., Jaquet, E., Winkler, C., & Clifford, C. W. (2004). Orientation-contingent face aftereffects and implications for face-coding mechanisms. *Current Biology, 14,* 2119–2123.

Rossion, B. (2008). Constraining the cortical face network by neuroimaging studies of acquired prosopagnosia. *Neuroimage, 40,* 423–426.

Rossion, B., Caldara, R., Seghier, M., Schuller, A. M., Lazeyras, F., & Mayer, E. (2003). A network of occipito-temporal face-sensitive areas besides the right middle fusiform gyrus is necessary for normal face processing. *Brain, 126,* 2381–2395.

Rotshtein, P., Vuilleumier, P., Winston, J., Driver, J., & Dolan, R. (2007). Distinct and convergent visual processing of high and low spatial frequency information in faces. *Cerebral Cortex, 17,* 2713–2724.

Schiltz, C., & Rossion, B. (2006). Faces are represented holistically in the human occipito-temporal cortex. *Neuroimage, 32,* 1385–1394.

Sergent, J., & Signoret, J. L. (1992). Varieties of functional deficits in prosopagnosia. *Cerebral Cortex, 2,* 375–388.

Steeves, J. K., Culham, J. C., Duchaine, B. C., Pratesi, C. C., Valyear, K. F., Schindler, I., et al. (2006). The fusiform face area is not sufficient for face recognition: Evidence from a patient with dense prosopagnosia and no occipital face area. *Neuropsychologia, 44,* 594–609.

Tong, F., Nakayama, K., Moscovitch, M., Weinrib, O., & Kanwisher, N. (2000). Response properties of the human fusiform face area. *Cognitive Neuropsychology, 17,* 257–279.

Tsao, D. Y., Freiwald, W. A., Tootell, R. B., & Livingstone, M. S. (2006). A cortical region consisting entirely of face-selective cells. *Science, 311,* 670–674.

Watson, T. L., & Clifford, C. W. (2006). Orientation dependence of the orientation-contingent face aftereffect. *Vision Research, 46,* 3422–3429.

Wilkinson, F., James, T. W., Wilson, H. R., Gati, J. S., Menon, R. S., & Goodale, M. A. (2000). An fMRI study of the selective activation of human extrastriate form vision areas by radial and concentric gratings. *Current Biology, 10,* 1455–1458.

Wilkinson, F., Wilson, H. R., & Habak, C. (1998). Detection and recognition of radial frequency patterns. *Vision Research, 38,* 3555–3568.

Wilson, H. R., Loffler, G., & Wilkinson, F. (2002). Synthetic faces, face cubes, and the geometry of face space. *Vision Research, 42,* 2909–2923.

Winston, J. S., Henson, R. N., Fine-Goulden, M. R., & Dolan, R. J. (2004). fMRI-adaptation reveals dissociable neural representations of identity and expression in face perception. *Journal of Neurophysiology, 92,* 1830–1839.

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.

Young, A. W., Hay, D. C., McWeeny, K. H., Flude, B. M., & Ellis, A. W. (1985). Matching familiar and unfamiliar faces on internal and external features. *Perception, 14,* 737–746.

Yovel, G., & Kanwisher, N. (2004). Face perception: Domain specific, not process specific. *Neuron, 44,* 889–898.

Yovel, G., & Kanwisher, N. (2005). The neural basis of the behavioral face-inversion effect. *Current Biology, 15,* 2256–2262.