

Selectivity for the configural cues that identify the gender, ethnicity, and identity of faces in human cortex

Minna Ng, Vivian M. Ciaramitaro, Stuart Anstis, Geoffrey M. Boynton, and Ione Fine

PNAS published online Dec 12, 2006;

doi:10.1073/pnas.0605358104

This information is current as of December 2006.

E-mail Alerts	This article has been cited by other articles: www.pnas.org#otherarticles
Rights & Permissions	Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or click here .
Reprints	To reproduce this article in part (figures, tables) or in entirety, see: www.pnas.org/misc/rightperm.shtml
	To order reprints, see: www.pnas.org/misc/reprints.shtml

Notes:

Selectivity for the configural cues that identify the gender, ethnicity, and identity of faces in human cortex

Minna Ng^{*†‡}, Vivian M. Ciaramitaro^{*†}, Stuart Anstis[†], Geoffrey M. Boynton^{*}, and Ione Fine[§]

^{*}The Salk Institute, 10010 North Torrey Pines Road, La Jolla, CA 92037; [†]Department of Psychology, University of California at San Diego, La Jolla, CA 92093; and [§]Department of Ophthalmology and Zilkha Neurogenetic Institute, University of Southern California, Los Angeles, CA 90033

Edited by Larry R. Squire, University of California, San Diego, CA, and approved October 31, 2006 (received for review June 27, 2006)

We used psychophysical and functional MRI (fMRI) adaptation to examine how and where the visual configural cues underlying identification of facial ethnicity, gender, and identity are processed. We found that the cortical regions showing selectivity to these cues are distributed widely across the inferior occipital cortex, fusiform areas, and the cingulate gyrus. These regions were not colocalized with areas activated by traditional face area localizer scans. Traditional face area localizer scans isolate regions defined by stronger fMRI responses to a random series of face images than to a series of non-face images. Because these scans present a random assortment of face images, they presumably produce the strongest responses within regions containing neurons that are face-sensitive but not highly tuned for face type. These areas might be expected to show only weak selective adaptation effects. In contrast, the largest responses to our selective adaptation paradigm would be expected within areas containing more selectively tuned neurons that might be expected to show only a sparse collective response to a series of random faces. Many aspects of face processing (e.g., prosopagnosia, recognition, and configural vs. featural processing) are likely to rely heavily on regions containing high proportions of neurons that show selective tuning for faces.

adaptation | functional MRI | cingulate gyrus | fusiform

When looking at an unfamiliar face, its gender and ethnicity tend to be more salient than almost any other feature. For example, when preparing “mug books” for identification of suspects, a standard law enforcement guide advocates grouping photos of individuals that are similar in ethnicity, age, and gender, because eye-witnesses rely strongly on these facial properties (1). Similarly, for preschool children, gender is the most salient characteristic for face categorization; ethnicity and age are also salient, but the presence or absence of eyeglasses is extremely nonsalient (2).

Recently, several groups have begun using psychophysical adaptation paradigms to examine sensitivity to the facial properties that underlie ethnicity, gender (3, 4), and identity (5) in humans. For example, Webster *et al.* (3) recently showed that adaptation to a set of female faces biases the perception of a face that is normally gender neutral (seen as male half the time and female half the time), making it appear male. This adaptation effect is analogous to the waterfall illusion, in which adaptation to upward motion makes a stationary pattern appear to move downward (6, 7). These adaptation effects seem not to be entirely mediated by low-level mechanisms, because adaptation effects can transfer across both size and retinal position (4, 8). A similar paradigm has also been used to demonstrate the existence of mechanisms tuned for the identity of particular faces (5). Adaptation to an individual face shifts the apparent perception of subsequently presented faces along a trajectory passing through the adapting face and an “average” face. An analogous functional MRI (fMRI) experiment has demonstrated that the same neural population responds to faces whose features are morphed away from a prototypical face along a single featural axis (9).

Numerous clinical (10, 11) and neuroimaging (12–16) studies have associated face processing with specialized regions in the human occipital and ventral temporal extrastriate cortex. Electrophysiology recordings on the surface of human inferior extrastriate visual cortex (17, 18) and analogous regions of the nonhuman primate inferior temporal cortex find neuronal selectivity for faces (19), including selectivity for particular face properties, such as identity (20), expression (21), and gender (22–24). These neurons are mostly invariant to low-level properties, such as size, position and contrast (19, 25). In recent years, there has been increasing interest in investigating the selectivity of neural tuning within face processing areas (9, 26, 27), as discussed below.

Here, we use psychophysical (3, 5) and fMRI adaptation (28, 29) to further examine how and where conjunctions of configural cues to ethnicity, gender and identity are processed in the human visual system. We were interested in whether cues to ethnicity, gender and identity were processed in similar regions of visual cortex, and to what extent these regions fall within cortical areas traditionally associated with face processing.

Results

Psychophysics. The psychophysical experimental design is shown in Fig. 1*a*. Subjects were presented with adaptor faces (for example, male Asian and female Caucasian face images) for an initial 3-min adaptation period. Face images were presented at a rate of one face per second. An auditory beep and a change in the size of the fixation point indicated the onset of each test trial. In each test trial, we presented a test image that was a morph either between male/female, or between Asian/Caucasian. (Morphing was carried out only along a single dimension, either gender or ethnicity, but not both.) The task of the subject was to judge the gender (male or female), or ethnicity (Asian or Caucasian) of the test image by means of a key-press. No feedback was given. Between every trial there was a 12-s top-up adaptation period.

If there exist jointly tuned mechanisms (Fig. 1*b Left*) which are selective for both gender and ethnicity cues, then adaptation to a combination of male Asian and female Caucasian faces should result in male faces appearing more Caucasian, and female faces appearing more Asian. Without this joint tuning, we would not expect a net adaptation effect (Fig. 1*b Right*) because male, female, Asian and Caucasian faces were presented with equal frequency

Author contributions: M.N., S.A., G.M.B., and I.F. designed research; M.N., V.M.C., and I.F. performed research; I.F. contributed new reagents/analytic tools; M.N., G.M.B., and I.F. analyzed data; and M.N., V.M.C., S.A., G.M.B., and I.F. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS direct submission.

Abbreviations: fMRI, functional MRI; MA, male Asian; FA, female Asian; MC, male Caucasian; FC, female Caucasian; BOLD, blood oxygenation level-dependent; iOcc, inferior occipital cortex; FuG, fusiform gyrus; CG, cingulate gyrus.

[†]To whom correspondence should be addressed. E-mail: mng@salk.edu.

This article contains supporting information online at www.pnas.org/cgi/content/full/0605358104/DC1.

© 2006 by The National Academy of Sciences of the USA

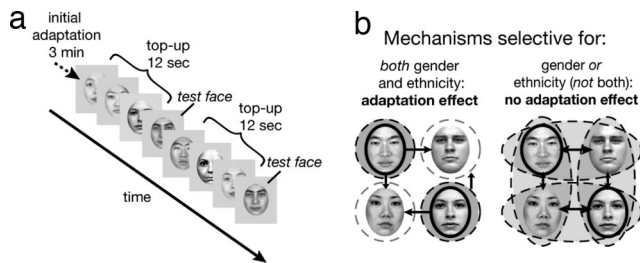


Fig. 1. Psychophysical experimental design and predictions. (a) Subjects were adapted to a series of images that randomly alternated between MA and FC faces (shown with black solid circles) and made ethnicity discriminations on test images morphed between MA and MC, or FA and FC faces. (b) *Left*) Dashed circles represent the mechanisms selective for both gender and ethnicity. *Right*) Dashed ovals represent the mechanisms selective for either gender or ethnicity. Mechanisms predicted to be adapted by MA and FC faces are shaded in gray.

during the adaptation period. Consequently, mechanisms specifically tuned for gender or ethnicity, but not for both, would be expected to adapt uniformly, resulting in no net change in the apparent ethnicity or gender of test faces. Analogous contingent after-effects have been used extensively to examine the tuning of low-level mechanisms. For example, in the McCollough effect, adaptation to an alternating green-black horizontal grating and red-black vertical grating produces an orientation-specific adaptation effect causing horizontal gratings to appear redder and vertical gratings to appear greener (30–35).

Fig. 2a shows example psychophysical data in which the subject judged the ethnicity of test images morphed between male Asian (MA) and male Caucasian (MC) both before and after adaptation to MA and female Caucasian (FC) faces. After adaptation, the male test images appear less Asian. Because the morph continuum is nonlinear, this adaptation effect was quantified along the y axis, as described in Fig. 2a. We ran all possible combinations of adapting and test stimuli in a counterbalanced design.

Significant contingent adaptation effects (ANOVA, $P < 0.05$) were found in every subject (Fig. 2b), demonstrating the efficacy of our paradigm in shifting the appearance of faces. The mean adaptation effect across subjects was 18.4%. This significant

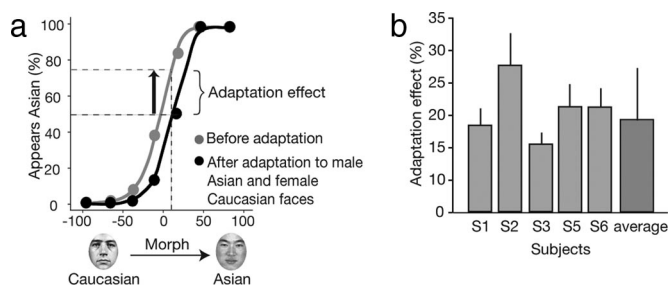


Fig. 2. Psychophysical results. (a) Example of psychometric functions before and after adaptation to FC and MA faces. The x axis represents the morph continuum and the y axis represents the percentage of time that the subject (S2) responded that the test face appeared Asian. Data points were fit by using a cumulative normal function. Because the x axis is not linear, we measured the shift in the psychometric functions due to adaptation along the y axis. We interpolated to find the morph that would be seen as Asian on 50% of the trials after adaptation (Morph no. 4, black dotted line). We then interpolated to the preadaptation psychometric function (gray dotted line) to find the percentage of trials that Morph no. 4 would be seen as Asian before adaptation (79%); this example yielded a 29% adaptation effect. (b) Psychophysical adaptation effects for all subjects. Error bars for individual subjects are calculated across all repeats of all conditions. The error bar for the mean response is calculated across subjects. (S1, 15.61%, SE \pm 2.85; S2, 26.52%, SE \pm 5.64; S3, 12.13%, SE \pm 2.04; S5, 18.9%, SE \pm 3.95; S6, 18.91%, SE \pm 3.25; mean, 18.41, SE \pm 8.23).

contingent adaptation effect suggests the presence of mechanisms selective for both gender and ethnicity (Fig. 1b).

fMRI. Face selective adaptation responses were measured in three conditions. In all adaptation scans, before blood oxygenation level-dependent (BOLD) signal measurement, subjects were preadapted for 3 min to an adaptor face set while lying in the scanner. Immediately after preadaptation, we measured BOLD responses using an uneven block design that alternated between presenting the adaptor set for 24 s, and the nonadapted set for 8 s.

The contingent adaptation condition was analogous to the psychophysical experiment described above. The adaptor set contained MA and FC faces; the nonadapted set contained female Asian (FA) and MC faces, or vice versa. In the individual adaptation condition, we arbitrarily assigned faces to the adaptor or nonadapted sets. In the configural adaptation condition, we adapted subjects to Asian and Caucasian hermaphrodites (morphs between males and females), and the nonadapted set consisted of Eurasian males and females (morphs between Asian and Caucasian), or vice versa.

We also carried out three traditional face area localizer conditions: (i) faces vs. houses, (ii) faces vs. cars, and (iii) faces vs. phase-scrambled faces, using a standard even block design (16 s on, 16 s off).

We used two predicted time courses to fit BOLD responses over time. The slow time course assumes that adaptation occurs on a relatively long time scale (several seconds or longer), and consequently in the case of adaptation scans, the neural response is dominated by the effects of the 4-min preadaptation period and the uneven duty cycle. The slow time course, therefore, predicts larger BOLD responses to nonadapted as compared with adapted stimuli. BOLD responses to localizer conditions were fit by using a standard time course that assumes that larger BOLD responses will be found for one of the two sets of stimuli (e.g., faces or houses).

The transient time course assumes that the majority of adaptation occurs on a much more rapid time scale (a few seconds or less), and therefore predicts a transient increase in BOLD response whenever the stimulus type changes, for both adaptation and localizer conditions. Supporting information (SI) Fig. 5 shows example BOLD responses that are better described by either a slow or transient time course.

Overall, fitting responses with a slow time course produced significantly more activation (number of voxels active at $P < 0.05$) than fitting with a transient time course ($P < 0.001$; three-factor ANOVA, subject \times condition \times time course). This analysis suggests that a significant proportion of the adaptation in our study occurred on a time scale of several seconds or longer. Further analyses were therefore based on slow time course predictors.

Fig. 3 shows parameter maps of coherence values based on the slow time course. Activity was distributed across inferior occipital cortex (iOcc) and fusiform gyrus (FuG) areas, as well as the cingulate gyrus (CG). Coherence maps based on the transient time course showed significantly less activity. However, where activation was found with the transient time course, the location of activity was qualitatively similar to the adaptation found by using the slow time course (see SI Fig. 6).

Across all four subjects we tend to see robust activity in phase with the nonadapted face images (red/magenta) for all three adaptation conditions. As might be expected, we see little or no out-of-phase activity (green) in any of the adaptation experiments. The other two localizers (data not shown) had similar patterns of activity as the face vs. house localizer condition.

In all our adaptation sessions, we used a finite stimulus set. Consequently, in a single scanning session, subjects were exposed to a given face as many as 29 times when it was from the adapting set and as few as 2 times when it was a face from the nonadapted set. Therefore, our contingent and configural conditions pre-

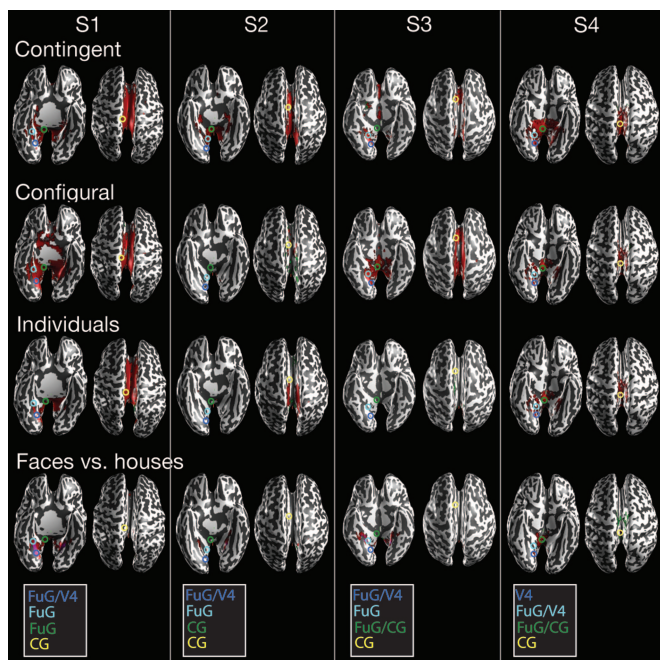


Fig. 3. Coherence maps based on the slow time course. Data were spatially blurred in the 3D representation by using a Gaussian filter with a half-width at half height of 6.7 mm. Responses in phase (using the slow time course) with the nonadapted face set (adaptation conditions) or face images (localizer) are in red ($P = 0.05$) and magenta ($P = 0.01$); activity in phase with the adapted face set or house images is in green ($P = 0.05$). P values of 0.05 and 0.01 correspond to coherence values 0.169 and 0.236, respectively. Areas V4, FuG, and CG are circled.

sumably include significant adaptation to individual faces, as well as adaptation to gender and ethnicity cues.

Averaged across subjects, the contingent and configural condition had 2.8 and 2.7 times as many active voxels, respectively, as the individual condition. However, a single-factor ANOVA comparing the number of voxels active at a coherence threshold of $P < 0.05$ (unblurred) across the three adaptation conditions fell below significance ($P = 0.07$). A 1-tailed paired t test ($P < 0.05$, Bonferroni adjusted) did find fewer voxels active in the individual adaptation condition than in the contingent adaptation condition. There was no significant difference in the number of active voxels between configural and individual adaptation conditions. Nor was there a significant difference in the number of active voxels between contingent and configural conditions, suggesting that neurons do not have a bias toward being tuned along male–female and Asian–Caucasian “cardinal axes” (36).

Activated areas were defined by using a combination of Talaraich coordinates (37) and anatomical landmarks (Table 1). As would be expected, we found areas within iOcc and FuG that responded selectively to face vs. non-face stimuli in the localizer scans (38).

In adaptation conditions, we see activity close to the regions that showed activation in response to faces in localizer conditions. However, in adaptation scans we also see activity in the CG. In the contingent condition, all subjects showed significant activity within the CG; in the individual and configural conditions, three of the four subjects showed significant activity within the CG. Very little activity was found within the CG for localizer conditions. Our finding of face-selective responses with CG may have implications for the difficulties in face processing found in autism spectrum disorders (39).

The blurring carried out for Fig. 3 requires that voxels over a 6-mm region or greater show a mean activity of at least $P = 0.05$. Such clustering criteria helps to reduce false positives due to

Table 1. Areas identified as showing significant activity for each subject

Condition	S1	S2	S3	S4
Contingent		iOcc	iOcc	iOcc
	FuG	FuG	FuG	FuG
	CG	CG	CG	CG
Configural	iOcc		iOcc	iOcc
	FuG		FuG	FuG
	CG	CG	CG	CG
Individual			iOcc	iOcc
	FuG		FuG	FuG
	CG	CG		CG
Faces vs. houses	iOcc		iOcc	iOcc
	FuG	FuG	FuG	FuG

multiple comparisons (40). However, one consequence of such clustering is that the activity is represented on a relatively coarse spatial scale.

It has recently been shown that individual voxels can show orientation selective responses that are driven by differences in selectivity across cortical distances of less than a millimeter (41, 42). However, these classification algorithms rely on analyzing responses within a predetermined region of interest (ROI). Moreover, these algorithms are susceptible to small differences in mean response between conditions, and to differences in the time course between the uneven and even block designs. Therefore, to quantify spatial overlap on a finer scale between conditions, we calculated the pair-wise cross-correlation for each voxel between coherence values (unblurred) in every possible pair of conditions. This analysis was done by using an ROI for each subject that included all voxels that showed significant activity, either in or out of phase, with the presentation of faces (blurred, $P < 0.05$) in all adaptation and localizer conditions.

Fig. 4a shows the mean correlation coefficient between each condition across all voxels averaged across subjects. A positive correlation coefficient (shown in red) indicates overlap between voxels of high coherence values between two conditions, i.e., both conditions produce similar spatial patterns of responses; negative correlations are shown in green. Fig. 4b and c shows the same cross-correlation analysis for activity within ventral and dorsal areas, respectively (14).

All three localizer conditions show strong mutual correlations that are presumably being driven by colocalized responses to faces across the three localizers. This correlation may be particularly strong because all three localizers were carried out within the same session, thereby minimizing differences in field distortions and misalignment across localizer scans. Responses to the contingent, configural and identity conditions were also mutually correlated, despite being run in separate sessions. This result suggests that similar areas were activated in the three adaptation conditions.

Surprisingly, responses in adaptation conditions are not particularly well correlated with responses in phase with faces in the localizer scans. Regions showing selective adaptation effects show surprisingly little spatial overlap with traditionally defined face areas. Indeed, correlations in activity between adaptation and localizer conditions were as likely to be negatively as they were likely to be positively correlated. Thus, regions showing adaptation to individuals or to gender and ethnicity show surprisingly little spatial overlap with traditionally defined face areas. This pattern of results was found in ventral as well as dorsal areas, suggesting that differences in spatial distribution exist on a relatively fine scale,

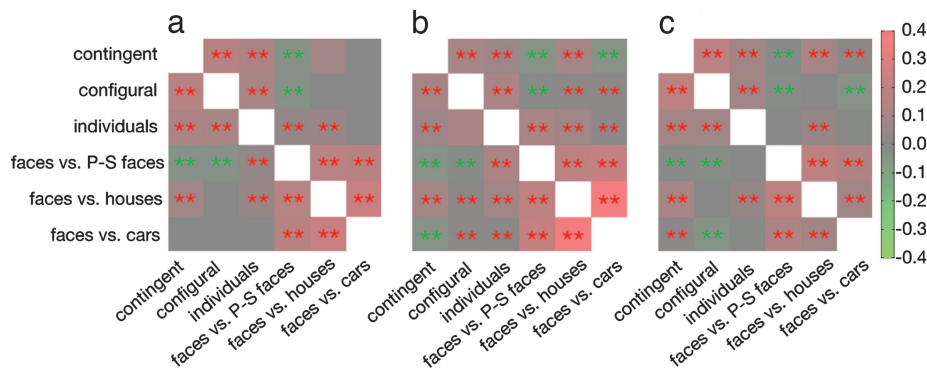


Fig. 4. Correlation coefficients between conditions, averaged across all voxels (a), ventral regions only (includes iOcc and FuG) (b), and dorsal regions only (includes CG) (c). Significance values of $P < 0.05$ (*) and $P < 0.01$ (**) were determined by using a Monte Carlo simulation procedure in which we scrambled the location of voxels. Because of the large number of voxels included in the analysis, even very low correlation values can be significant. Significantly positively and negatively correlated activity is indicated by red and green asterisks, respectively.

consistent with recent findings of heterogeneity of selectivity within ventral regions using high-resolution imaging (43).

Discussion

Face processing involves many stages: first, we recognize that a face is present; then we further classify the face or identify it as an individual that we know. It has previously been shown that within areas isolated by traditional face area localizers, there are subdivisions that differentiate between different types of face processing. For example, within the areas that respond to face localizer scans, the iOcc tends to be sensitive to the physical information present in a face, whereas the right FuG tends to be more sensitive to identity information (26). Another recent study has shown that the FuG and posterior superior temporal sulcus tend to be more sensitive to identity information, whereas information about emotion tends to be processed in a more anterior portion of the superior temporal sulcus (27). Previous studies have also found evidence for distributed networks for face processing that can fall outside traditionally defined areas (27, 44).

Here, we examined sensitivity to the configural cues that underlie face processing. It has previously been shown psychophysically that subjects adapt to gender, and ethnicity cues (3), as well as to individual identity (5). Here, using a slow adaptation technique, we measured adaptation to conjunctions of configural cues using both psychophysical and fMRI adaptation.

The slow adaptation paradigm used in our fMRI experiment differs significantly from those used in previous rapid event-related studies, which measure responses to a rapid change of identity (or a change along some other configural dimension) as compared with the immediate repetition of a single face (9, 26, 27). Worryingly, it has been shown that fMRI measurements of selectivity can differ depending on whether fast or slow adaptation paradigms are used. For example, in V1 orientation-selective release from adaptation is only observed using a slow adaptation technique (45). In our study, we examined both fast and slow adaptation effects by fitting data with both transient and slow predicted time courses. We found that slow adaptation effects (over several seconds) were significantly more powerful than transient adaptation effects (occurring over a few seconds). However, where activation was found with the transient time course, the location of activity was qualitatively similar to the adaptation found on a slow time scale (see Fig. 3 and SI Fig. 6).

We find here, that the regions of cortex showing selectivity for the configural cues that identify gender, ethnicity, and identity fall within the iOcc, FuG, and CG. There was no significant difference in the extent of activation between the configural and the contingent condition, suggesting a lack of strong biases in tuning toward cardinal axes of male–female and Asian–Caucasian. Activation was

slightly stronger in configural and contingent adaptation conditions than in the individual adaptation condition, but this effect was below significance. There was no difference in the magnitude of adaptation effects between configural and contingent conditions, suggesting that adaptation is driven by configural similarities along any orientation, rather than clustering along “cardinal” dimensions of gender and ethnicity.

We expected to find adaptation effects in the FuG, because activity in response to faces is regularly observed in this area. However, the robust adaptation effects in the CG were less expected. This area has been associated much more consistently with a variety of other tasks, including error processing, selective and competitive attention, expectancy and reward (46–48). One possibility is that responses within the CG were driven by the greater “novelty” of the nonadapted faces (or some other state-change such as expectancy that might differ across even and uneven block designs). However, if this possibility were the case we might have expected to find responses within the CG to the transient time course for both adaptation and localizer conditions (assuming that the CG was sensitive to novelty changes on a relatively short time scale of a few seconds). There was no evidence supporting this possibility in our data, as shown in SI Fig. 6.

One previous study (49) has found the CG to be associated with the encoding of a variety of faces, and a second (50) found CG activity during a reverse learning paradigm that involved learning and responding to contingencies between face identity and emotional expression. Our data suggest that highly selective face sensitive neurons do exist within the CG. It is not entirely surprising that this area has seemed unresponsive in many previous studies: the traditional localizer approach is not well suited for identifying areas in which neurons are face-sensitive and highly selective in their tuning. In areas containing neurons with highly selective tuning, any given neuron would rarely (if ever) be activated when presented with a series of random faces, resulting in a weak collective fMRI response to traditional localizer scans.

Our cross-correlation technique enables examination of the spatial distribution of activity patterns on a fine-scale. Such patterns would be masked by blurring, clustering, or averaging of the fMRI response across different observers’ brains (26, 27). We found that adaptation responses across contingent, configural and identity conditions were correlated with each other, suggesting similar spatial distributions of activity. Overall, these results are consistent with the notion that the dimensions of gender, ethnicity and identity are simply different dimensions along which facial similarity can be varied, without necessarily involving separate processing regions.

Regions showing adaptation effects were relatively uncorrelated with regions responding to faces in the localizer scans, within both

dorsal and ventral regions of cortex. The lack of spatial correlation between traditionally defined face regions and adaptation-sensitive regions is likely to be due to the insensitivity of traditional localizer scans to produce responses within areas containing highly selective neurons. Because traditional localizers present a random assortment of faces, strong responses to localizer scans will be found within areas containing very broadly tuned face-sensitive neurons. However, neurons that are very broadly tuned will respond to all faces regardless of their ethnicity or gender, and will therefore show weak or nonexistent selective adaptation effects. In contrast, as described above, very selectively tuned neurons will only show a weak collective response to a series of random faces. For many studies examining face processing (e.g., prosopagnosia, recognition, configural vs. featural processing, and the face processing deficits associated with autism spectrum disorders), regions showing selective face tuning are likely to be at least as important as those less selective regions that tend to be identified by using traditional localizer scans.

Materials and Methods

Psychophysics. Participants. Six subjects (four females and two males; 23–29 years of age; mean age, 24.6) with normal vision gave informed consent to participate in this experiment, which was approved by The Salk Institute of Biological Studies institutional review board. One subject did not show adaptation under any condition, and his data have been excluded.

Stimuli. Adapting faces were frontal-view gray-scale (256×256 pixel) images of Asian (A), Caucasian (C), male (M), and female (F) faces of neutral-expression. Some face images were from the Ekman 1976 face set and the Cohn–Kanade AU-Coded Facial Expression Database; others were photographs of students and staff of UCSD and The Salk Institute.

Pairs of adaptor faces were used to create test morph images (MorphMan, version 4.0; STOIK Imaging, Moscow, Russia) varying across either ethnicity or gender (but not both). We used a method of constant stimuli to present the morphed images. Seven faces, evenly distributed along the morph between 100% male and 100% female, were presented an equal number of times in a random order.

Each face image was presented on a computer monitor with a spatial resolution of $1,600 \times 1,200$ at a viewing distance of 57 cm, and each face image subtended $\approx 7.1^\circ$ of visual angle.

Procedure. Subjects were presented with adaptor faces (for example, MA and FC face images) for an initial 3-min adaptation period at a rate of one face per second (Fig. 1*a*). A beep and a change in the size of the fixation point indicated the onset of every test trial. In each test trial, we presented a test image that was a morph either between male/female, or between Asian/Caucasian. (Morphing was carried out only along a single dimension, either gender or ethnicity, but not both.) The subject judged the gender (male or female) or ethnicity (Asian or Caucasian) of the test image by means of a key-press. Between every test trial was a top-up adaptation period of 12 s. Each session contained ≈ 56 trials. In the baseline condition, subjects performed the same task, but there was no initial adaptation period, and a fixation point on a blank screen appeared during the top-up period.

Subjects carried out six to seven baseline (no adaptation) sessions (≈ 364 trials), and eight sessions of contingent adaptation, where they adapted to a conjunction of gender and ethnicity cues (448 trials). Two sessions were carried out on each testing day. On any given testing day the same adapting face categories were used throughout both testing sessions to maximize adaptation effects. For example, in session one, a subject might be adapted to MA and FC faces, and judge the gender of test images; in session two (same day), the subject would again be adapted to MA and FC faces, but would make an ethnicity judgement. On the next testing day, the subject would be adapted

to FA and MC faces. The order of testing sessions was randomly counterbalanced across subjects. The four testing days needed to carry out the eight contingent sessions were scheduled over a 3-week period. It should be noted that the effect of any residual adaptation from a previous day's testing would only serve to reduce the magnitude of measured adaptation.

MRI. Participants. Four subjects (two females and two males; 23–29 years of age; mean age, 25.3) gave informed consent to participate in this experiment, which was approved by the institutional review board of The Salk Institute for Biological Studies. Each subject participated in six scanning sessions. Three of these subjects (S1–S3) also participated in the psychophysics experiment. fMRI data were acquired several weeks after the completion of the psychophysical experiment.

Stimuli. Face images were the same as the unmorphed faces used for adaptation in the psychophysical experiment. There were 11 exemplars of each face type, and images were again presented at a rate of 1 s per image (similar to the psychophysical experiment). Object images included 22 exemplars each of cars, houses, and phase-scrambled faces. The images subtended $\approx 6.7^\circ$ of visual angle (as compared with 7.1° of visual angle in the psychophysical experiment).

Apparatus. Functional imaging was conducted at the Center for Functional Magnetic Resonance Imaging at the University of California, San Diego (UCSD). A GE 3T scanner system and standard eight-channel head coil was used to collect the functional data; a Siemens 1.5T scanner was used to collect an anatomical reference volume, which was used to align functional data across multiple scanning sessions by using standard alignment techniques (51). See *SI Methods*.

In all sessions, subjects lay in a supine position stabilized by a bite bar, which was fastened to the magnetic resonance table. In functional scans, visual stimuli were projected from a computer laptop by an NEC projector onto a screen near the subject's head. Subjects viewed the reflected stimuli on an angled mirror fixed above the eyes.

Adaptation scans. Only one scanning session (one condition) was conducted in a day, to avoid subject fatigue and minimize carry-over adaptation across sessions. The effect of any residual adaptation remaining from adaptation on a previous day would simply serve to reduce the magnitude of measured adaptation.

Before BOLD signal measurement, subjects were preadapted for 4 min to faces from the adaptor set, while lying in the scanner. We then immediately measured BOLD responses using an uneven block design, which alternated between the adaptor set (24 s) and the nonadapted set (8 s). This procedure was repeated for six and a half cycles for a total scan time of 208 s. The first half cycle was discarded to avoid magnetic saturation effects; thus, analyzed data consisted of six cycles. The transition between preadaptation and measurement was invisible to the observer, except for a change in scanner noise.

Contingent adaptation. We carried out two sessions, each consisting of six repeated scans, to measure contingent adaptation to a combination of gender and ethnicity cues. In one session, the adaptor set consisted of MA and FC faces, and the nonadapted set consisted of FA and MC faces. In the second session (carried out on a different day), the faces in the adaptor set would consist of MC and FA faces, and the nonadapted faces would consist of MA and FC faces. The order of these sessions was randomly counterbalanced across subjects, and the order in which face images were presented was randomized both within and across scans. In total, we obtained 48 contingent condition scans (two sessions \times six repetitions \times four subjects).

Individual adaptation. To measure adaptation to individual faces, we arbitrarily assigned each face to either the adaptor or nonadaptor set. The same adapting set was used throughout a given session to maximize adaptation effects. This condition was equiv-

alent to the identity specific adaptation demonstrated psychophysically by Leopold *et al.* (ref. 5; also see ref. 9). In total, we obtained 24 individual condition scans (six repetitions \times four subjects).

Configural adaptation. To test the extent to which any adaptation observed in the contingent condition was due to general configural similarities within faces, we adapted three of our four subjects to Asian and Caucasian hermaphrodites (morphs between males and females) and measured adaptation effects for Eurasian males and females (morphs between Asian and Caucasian), and vice versa. If mechanisms are preferentially tuned along the cardinal directions of M/F, and A/C, then, analogous to experiments that have been carried out examining chromatic tuning (52), we would expect to see weaker effects when adapting “off-axis.”

We carried out two scanning sessions, each consisting of six repeated scans. In one session, the adaptor set contained male Eurasian and female Eurasian faces, and the nonadapted set consisted of hermaphrodite Caucasian and hermaphrodite Asian faces. In the second session (carried out on a different day), the adaptor set was hermaphrodite Caucasian and Asian face images, and the nonadapted set consisted of male and female Eurasian faces. The order of sessions was randomly counterbalanced across subjects. In total, we obtained 48 scans for the configural condition (two sessions \times six repetitions \times four subjects).

Faces vs. non-face localizer scans. We ran the following three localizers: (i) faces vs. houses, (ii) faces vs. cars, and (iii) faces vs. phase-scrambled faces. We used an even-block design that alternated between a variety of faces for 16 s, and, for example, a variety of car images for 16 s; this cycle was repeated six and a half times, for a total of 208 s of scan time. Analyzed data consisted of six cycles because the first half cycle was discarded to avoid magnetic saturation effects. Each localizer was run twice for a total of 12 scans per localizer condition. All six localizer scans were completed in a single session. In total, we obtained

24 localizer scans (two repetitions \times three localizer conditions \times four subjects).

Analysis. Segmentation, flattening, and inflation of cortical surfaces were carried out by using customized Matlab software (53, 54). Linear trends were subtracted from the fMRI time course of each voxel and the activity of each voxel across each scan was divided by the mean activity to convert BOLD response to percentage signal change (55). There were, necessarily, slight differences in slice prescriptions across scans. Occipital and temporal areas were always included, but the extent of parietal coverage varied slightly across sessions. Each session’s data were registered to an anatomical image collected in a separate scanning session, and for each subject only those voxels for which we collected data in every session were included in further analyses.

The fMRI response of each voxel was then fit with an estimate of the hemodynamic response function. In the case of the adaptation scans, a canonical hemodynamic impulse response function was convolved with the uneven block-design (51), and coherence values were calculated between the response of each voxel and the estimated BOLD responses for the best fitting delay.

In the case of the localizer scans, the impulse response function was convolved with the even block design, and coherence values were calculated between the response of each voxel and the estimated BOLD responses for the best fitting delay for responses in phase with the face stimuli. Voxels that responded out of phase with the non-face stimuli were assigned a negative coherence value. We then averaged coherence and delay values for each voxel across every repetition of each condition.

We thank M. Webster (University of Nevada, Reno, NV) for help and advice. This work was supported by National Science Foundation Integrative Graduate Education and Research Traineeship Fellowships DGE-0333451 (to M.N.), NEI-012925 (to G.M.B.), and NEI-014645 (to I.F.).

- Department of Justice, Office of Justice Programs (1999) *Eyewitness Evidence: A Guide for Law Enforcement* (Department of Justice, National Institute of Justice, Washington, DC).
- McGraw KO, Durm MW, Durnam MR (1989) *J Genet Psychol* 150:251–267.
- Webster MA, Kaping D, Mizokami Y, Duhamel P (2004) *Nature* 428:557–561.
- Webster MA, MacLin OH (1999) *Psychon Bull Rev* 6:647–653.
- Leopold DA, O’Toole AJ, Vetter T, Blanz V (2001) *Nat Neurosci* 4:89–94.
- Movshon JA, Lennie P (1979) *Nature* 278:850–852.
- Mather G, Verstraten F, Anstis S (1998) *The Motion Aftereffect: A Modern Perspective* (MIT Press, Cambridge, MA).
- Zhao L, Chubb C (2001) *Vision Res* 41:2979–2994.
- Loffler G, Yourganov G, Wilkinson F, Wilson HR (2005) *Nat Neurosci* 8:1386–1390.
- Damasio AR, Damasio H, Van Hoesen GW (1982) *Neurology* 32:331–341.
- Farah MJ, Levinson KL, Klein KL (1995) *Neuropsychologia* 33:661–674.
- Sergent J, Signoret JL (1992) *Philos Trans R Soc London B Biol Sci* 335:55–61.
- Sergent J, Ohta S, MacDonald B (1992) *Brain* 115:15–36.
- Puce A, Allison T, Gore JC, McCarthy G (1995) *J Neurophysiol* 74:1192–1199.
- Kanwisher N, McDermott J, Chun MM (1997) *J Neurosci* 17:4302–4311.
- Grill-Spector K, Knouf N, Kanwisher N (2004) *Nat Neurosci* 7:555–562.
- Fried SI, Munch TA, Werblin FS (2002) *Nature* 420:411–414.
- Allison T, Ginter H, McCarthy G, Nobre AC, Puce A, Luby M, Spencer DD (1994) *J Neurophysiol* 71:821–825.
- Perrett DI, Rolls ET, Caan W (1982) *Exp Brain Res* 47:329–342.
- Yamane S, Kaji S, Kawano K (1988) *Exp Brain Res* 73:209–214.
- Hasselmo ME, Rolls ET, Baylis GC (1989) *Behav Brain Res* 32:203–218.
- Desimone R, Albright TD, Gross CG, Bruce C (1984) *J Neurosci* 4:2051–2062.
- Fried I, Cameron KA, Yashar S, Fong R, Morrow JW (2002) *Cereb Cortex* 12:575–584.
- Baylis GC, Rolls ET, Leonard CM (1985) *Brain Res* 342:91–102.
- Rolls ET, Baylis GC (1986) *Exp Brain Res* 65:38–48.
- Rotshtein P, Henson RN, Treves A, Driver J, Dolan RJ (2005) *Nat Neurosci* 8:107–113.
- Winston JS, Henson RN, Fine-Goulden MR, Dolan RJ (2004) *J Neurophysiol* 92:1830–1839.
- Grill-Spector K, Malach R (2001) *Acta Psychol (Amst)* 107:293–321.
- Avidan G, Hasson U, Hendler T, Zohary E, Malach R (2002) *Curr Biol* 12:964–972.
- McCullough C (1965) *Science* 149:1115–1116.
- Stromeyer CF, 3rd, Dawson BM (1978) *Perception* 7:407–415.
- DeValois RL (1978) *Invest Ophthalmol Vis Sci* 17:834–835.
- Gibson JJ, Radner M (1937) *J Exp Psychol* 20:453–467.
- Lovegrove WJ, Over R (1972) *Science* 176:541–543.
- Wandell BA, Poirson AB, Newsome WT, Baseler HA, Boynton GM, Huk A, Gandhi S, Sharpe LT (1999) *Neuron* 24:901–909.
- Engel S, Zhang X, Wandell B (1997) *Nature* 388:68–71.
- Talairach J, Tournoux P (1988) *Co-Planar Stereotaxic Atlas of the Human Brain* (Thieme Medical Publishers, New York).
- Haxby JV, Hoffman EA, Gobbini MI (2000) *Trends Cognit Sci* 4:223–233.
- Schultz R, Robins D (2005) *Functional Neuroimaging Studies of Autism Spectrum Disorders. Handbook of Autism and Pervasive Development Disorders*, eds Volkmar FR, Paul R, Klin A, Cohen D (Wiley, Hoboken, NJ), 3rd ed, pp 515–533.
- Friston KJ, Rotshtein P, Geng JJ, Sterzer P, Henson RN (2006) *Neuroimage* 30:1077–1087.
- Kamitani Y, Tong F (2005) *Nat Neurosci* 8:679–685.
- Haynes JD, Rees G (2005) *Nat Neurosci* 8:686–691.
- Grill-Spector K, Sayres R, Ress D (2006) *Nat Neurosci* 9:1177–1185.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001) *Science* 293:2425–2430.
- Fang F, Murray SO, Kersten D, He S (2005) *J Neurophysiol* 94:4188–4195.
- Shidara M, Richmond BJ (2002) *Science* 296:1709–1711.
- Carter CS, Braver TS, Barch DM, Botvinick MM, Noll D, Cohen JD (1998) *Science* 280:747–749.
- Petit L, Courtney SM, Ungerleider LG, Haxby JV (1998) *J Neurosci* 18:9429–9437.
- Haxby JV, Ungerleider LG, Horowitz B, Maisog JM, Rapoport SI, Grady CL (1996) *Proc Natl Acad Sci USA* 93:922–927.
- Kringelbach ML, Rolls ET (2003) *NeuroImage* 20:1371–1383.
- Teo PC, Sapiro G, Wandell BA (1997) *IEEE Trans Med Imaging* 16:852–863.
- Krauskopf J, Williams DR, Heeley DW (1982) *Vision Res* 22:1123–1131.
- Engel SA, Glover GH, Wandell BA (1997) *Cereb Cortex* 7:181–192.
- Sereno MI, McDonald CT, Allman JM (1994) *Cereb Cortex* 4:601–620.
- Boynton GM, Engel SA, Glover GH, Heeger DJ (1996) *J Neurosci* 16:4207–4221.