

Attention Selectively Modifies the Representation of Individual Faces in the Human Brain

Caterina Gratton,¹ Kartik K. Sreenivasan,¹ Michael A. Silver,^{1,2} and Mark D'Esposito¹

¹Helen Wills Neuroscience Institute and ²School of Optometry, University of California, Berkeley, California 94720

Attention modifies neural tuning for low-level features, but it is unclear how attention influences tuning for complex stimuli. We investigated this question in humans using fMRI and face stimuli. Participants were shown six faces (F1–F6) along a morph continuum, and selectivity was quantified by constructing tuning curves for individual voxels. Face-selective voxels exhibited greater responses to their preferred face than to nonpreferred faces, particularly in posterior face areas. Anterior face areas instead displayed tuning for face categories: voxels in these areas preferred either the first (F1–F3) or second (F4–F6) half of the morph continuum. Next, we examined the effects of attention on voxel tuning by having subjects direct attention to one of the superimposed images of F1 and F6. We found that attention selectively enhanced responses in voxels preferring the attended face. Together, our results demonstrate that single voxels carry information about individual faces and that the nature of this information varies across cortical face areas. Additionally, we found that attention selectively enhances these representations. Our findings suggest that attention may act via a unitary principle of selective enhancement of responses to both simple and complex stimuli across multiple stages of the visual hierarchy.

Introduction

Attention powerfully modifies the way we see the world, allowing us to filter vast amounts of information and to preferentially process those components that are most relevant to our current behavioral goals. However, despite the complexity of natural visual scenes, most studies examining modulation of neural representations by attention have focused on simple stimuli with low-level features. In these studies, sustained top-down attention selectively modifies representations in early visual areas by enhancing responses to stimuli at attended locations or containing attended features (see, for example, Gandhi et al., 1999; McAdams and Maunsell, 1999; Treue and Maunsell, 1999; Serences et al., 2009). In contrast, less is known about the effects of attention on representations of more complex stimuli that are represented in higher visual areas. While attention can influence the processing of complex categories such as faces and scenes by enhancing responses across entire category-selective areas of cortex (Corbetta et al., 1991; Wojciulik et al., 1998; Gazzaley et al., 2005), the mechanism by which attention modulates the tuning of responses to individual stimuli within these categories is unknown.

Faces are extremely prevalent and salient stimuli in our environment. Individual neurons (Gross et al., 1972) (for review, see Gross, 2011) and entire brain regions (Kanwisher et

al., 1997; McCarthy et al., 1997) exhibit specialized responses to faces compared with other stimuli. These regions comprise a distributed network (Fairhall and Ishai, 2007; Moeller et al., 2008; Zhang et al., 2009) that includes the occipital face area (OFA), the fusiform face area (FFA), the face-selective portion of the superior temporal sulcus (fSTS), and more anterior parts of the inferior temporal lobe (Haxby et al., 2000). Both face features (Freiwald et al., 2009) and sets of face morphs (Leopold et al., 2006) are represented in a continuous fashion by neurons within these areas, perhaps within the context of an orderly face space (Valentine, 1991). Thus, because of their systematic representations and their high salience, face stimuli are ideal for examining whether the effects of attention on complex feature tuning mirror the effects that have previously been reported for low-level feature representations.

However, due to the relatively coarse spatial resolution of functional MRI (fMRI), in which the smallest volumes of measurement (voxels) are typically several cubic millimeters and contain thousands of neurons with different stimulus preferences, it has proven difficult to detect signals that are selective for individual faces. Nevertheless, recent evidence from both fMRI adaptation (fMR-A) paradigms (Loffler et al., 2005; Rotshtein et al., 2005; Gilaie-Dotan and Malach, 2007) and multivoxel pattern classifiers (MVPA) (Kriegeskorte et al., 2007; Natu et al., 2010; Nestor et al., 2011) suggests that it may be possible to detect tuning for individual faces within face-selective regions. Here, we adapted a method recently developed by Serences et al. (2009) for constructing feature tuning curves for single voxels and demonstrate that single voxels carry information about individual faces. In addition, we found that directing attention to one of a pair of superimposed faces selectively enhanced responses in voxels representing the attended face.

Received Aug. 27, 2012; revised Jan. 25, 2013; accepted March 5, 2013.

Author contributions: C.G., K.S., M.A.S., and M.D. designed research; C.G. performed research; C.G. analyzed data; C.G., K.S., M.A.S., and M.D. wrote the paper.

This work was supported by grants from the NIH (MH63901 and NS40813 to M.D.), the NSF (GRFP to C.G.), and the DOD (NDSEG to C.G.).

The authors declare no competing financial interests.

Correspondence should be addressed to Caterina Gratton, University of California, Berkeley, 132 Barker Hall, Berkeley, CA 94720. E-mail: cgratton@berkeley.edu.

DOI:10.1523/JNEUROSCI.4142-12.2013

Copyright © 2013 the authors 0270-6474/13/336979-11\$15.00/0

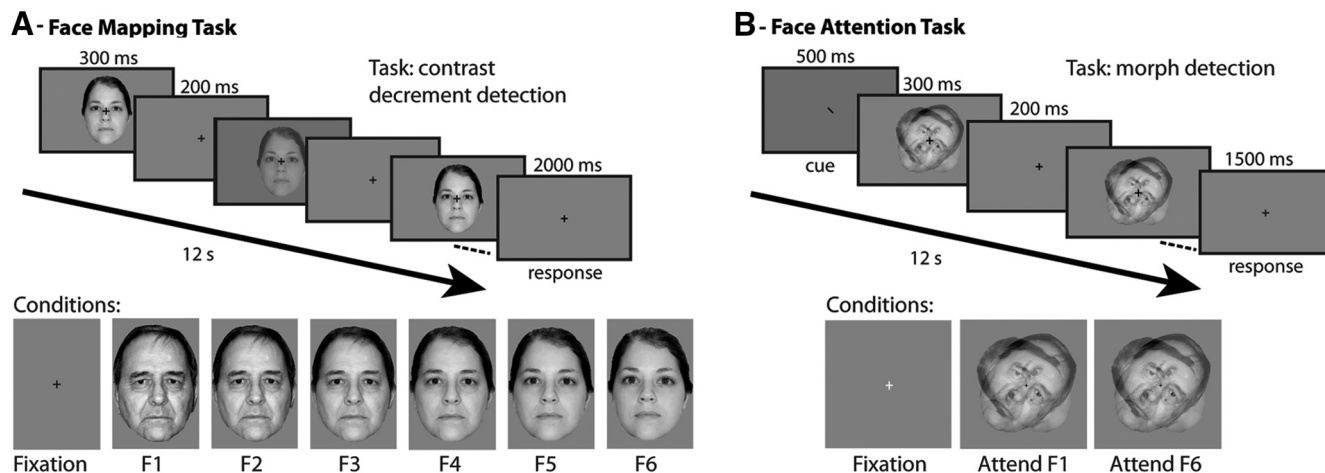


Figure 1. Trial schematic for face mapping (Experiment 1) and face attention (Experiment 2) tasks. **A**, Participants detected brief contrast decrements over the entire face during presentation of a series of flashing faces. Blocks were either fixation or one of six faces along a morph continuum (F1–F6). **B**, Attention blocks were either fixation or flashing superimposed images of F1 and F6. Participants attended to either F1 or F6 and detected brief morphs in the attended face toward a third face while ignoring morphs in the unattended face. The face to be attended was cued by the color (red or green) of the fixation cross. For illustrative purposes, the fixation crosses and contrast decrement displayed here are larger than those actually presented to participants, and fixation crosses are shown here in black.

Materials and Methods

Participants

Ten healthy participants (7 female, mean age = 25, age range 19–32) each completed two separate 2 h scanner sessions that were each accompanied by a 1 h practice session within the preceding week. Informed consent was obtained from subjects in accordance with procedures approved by the Committee for the Protection of Human Subjects at the University of California, Berkeley.

Behavioral tasks

Face localizer task. In each fMRI session, participants completed a face localizer run of 336 s either once or twice. This task consisted of sequential 16 s blocks of faces, scenes, and fixation. During face and scene blocks, participants performed a 1-back task, pressing a response button whenever an image was identical to the previous image. Faces and scenes were presented at a rate of one per second.

Face mapping task (Experiment 1). In each fMRI session, participants completed a 5 min face mapping run five or six times. A block consisted of either continuous presentation of a central fixation cross or repeated presentation of a single face, during which participants counted the number of contrast decrements in the face that occurred during the block. Each run consisted of 21 12 s blocks [3 blocks for each of seven conditions (6 face morphs and fixation), Fig. 1A]. In each block (except fixation), a single face was flashed on the screen for 10 s at a rate of 2 Hz (300 ms on, 200 ms off), followed by a 2 s response period during which the fixation cross was red. Fixation blocks consisted of 12 s of a white fixation cross. Block order was pseudo-randomized, with the constraint that the same condition was never repeated in two successive blocks. Participants were instructed to count the number of times that the contrast of the entire face was reduced for the duration of a single 300 ms presentation (0–3 contrast decrements possible per block, 25% probability of a given number of contrast decrements for each block). The amount of contrast decrement in the face was adjusted across runs to maintain ~75% correct trials (chance performance = 25%).

Stimuli consisted of grayscale images of an old male face (F1), a young female face (F6), and four morphs along a continuum between F1 and F6 (F2–F5; Fig. 1A). Images were obtained from the Center for Vital Longevity's face database (Minear and Park, 2004) and were morphed using Abrasoft's FantaMorph4 software after selecting a number of common image points in F1 and F6. The two faces on either end of the morph continuum were chosen to span a large range of facial features, including both gender and age. Intermediate faces were evenly spaced along the morph continuum (e.g., F2 was 80% F1 and 20% F6). After morphing, all faces were matched for average luminance and contrast (SD of the lumi-

nance) across the entire face and placed on a grayscale background, leaving in ears and hair. Face images were 12 degrees of visual angle along the longer vertical axis (forehead to chin).

Face attention task (Experiment 2). In each session, participants completed 6–7 runs of the 5 min attention task in addition to the face mapping task (Experiment 1) described above. Participants viewed a superimposed image of F1 and F6, with each face tilted 45 degrees in opposite directions, and were instructed to attend to just one of the two faces and to detect brief morphs in that face (Fig. 1B). Each run consisted of 24 12 s blocks, with 8 blocks in each of three conditions: attend F1, attend F6, or fixation. In attend F1 and attend F6 conditions, blocks began with a cue consisting of a red or green oriented bar indicating the identity of the face that was to be attended. The color associated with a particular face attention condition (e.g., green for attend F1) was consistent for a given subject but counterbalanced across subjects. The orientation of the cue indicated the orientation of the face to be attended. The cue remained on the screen for 500 ms, followed by the onset of a flashing composite image of F1 and F6 in the center of the screen (Fig. 1B). Face orientations were counterbalanced and pseudorandomly assigned across blocks in a scan. The composite face flashed at a rate of 2 Hz (300 ms on, 200 ms off) for 10 s. The fixation cross (red or green) remained on the screen throughout the block. Each block ended with a response period (1500 ms) during which the fixation cross was black.

The participants' task was to count the number of morphs in the attended face toward a third face (0–3 morphs per block, number of morphs randomly chosen and occurring at random times within the trial after the first flash). Morphs, when present, were displayed continuously for an entire flash (300 ms). The percentage morph to the third face was titrated for each subject and each attention condition (attend F1 or attend F6) to maintain ~75% correct trials (chance performance = 25%). The third face (i.e., 100% morph) was never shown to the subjects in any condition. The unattended face also morphed with the same probability as the attended face, but participants were instructed to ignore these changes. Block order was pseudorandom.

MRI procedures

MRI data acquisition. T2*-weighted echo planar images (EPI) were collected on a whole-body 3 Tesla Siemens MAGNETOM Trio MR scanner using a 32-channel radiofrequency head coil. Each localizer run consisted of 168 volumes (TR = 2 s, TE = 25 ms, FA = 70 deg, voxel size = 3 × 3 × 3.3 mm, with 34 slices ~12° from axial, no GRAPPA acceleration). Each face mapping run (Experiment 1) was 132 volumes (TR = 2 s, TE = 20 ms, FA = 70 deg, voxel size = 2.22 × 2.22 × 2.3 mm, with 46 slices ~12° relative to axial, GRAPPA factor = 4).

Each face attention run (Experiment 2) consisted of 150 volumes with the same acquisition parameters as the mapping runs. Due to scanner malfunction, face mapping runs were acquired with the localizer run parameters in three sessions, and face attention runs were acquired with the localizer run parameters in two sessions and with slightly altered parameters (matching typical face attention runs but with voxel size = $2.22 \times 2.22 \times 3.3$ mm, 34 slices) in one session. The acquisition volume covered approximately the whole brain, although it was positioned to cover all of anterior temporal cortex at the expense of some regions in the parietal cortex. Localizer runs had larger voxels to obtain a higher signal-to-noise ratio, whereas face mapping and attention runs had smaller voxel sizes to more densely sample the face regions for voxel tuning analyses. Structural images were acquired using a high-resolution MP-RAGE T1-weighted sequence.

MRI data analysis. Image preprocessing was conducted in AFNI [http://afni.nimh.nih.gov/afni/ (Cox, 1996)] and consisted of the removal of nonbrain voxels from the EPI volumes, despiking, motion correction (within and between runs of the same type), and rigid-body alignment to the structural image. In addition, spatial smoothing (4 mm FWHM) was applied to localizer runs.

General linear models (GLM) were computed for each type of run (localizer, face mapping, and face attention) to estimate response amplitude for each condition. All GLMs included head motion and detrending (third-degree polynomial) regressors. The GLM for the localizer task contained separate regressors for face and scene blocks. For the face mapping task (Experiment 1), separate GLMs were computed for each run (to generate independent β estimates for tuning analyses), and each GLM included regressors for each of the six face conditions and fixation, a response period regressor, and a single parametric “adaptation regressor.”

The motivation for including the adaptation regressor was to mitigate possible effects of face adaptation on our estimates of response amplitude. For example, because F1 and F6 are on the ends of the face morph continuum and therefore less likely than F2–F5 to be preceded by a stimulus nearby in face space, they would, on average, be the least adapted of the face stimuli. To account for this, we included a parametric regressor in our GLM estimation of the response weights that was scaled according to the relative distance, in face space, between the current block and the previous block. Thus, if the current block was very similar to the previous block in face space (e.g., the current block was F1 and the previous block was F2), this regressor would be assigned a value of 1, indicating the high potential contribution of face adaptation. If, instead, the current block was very different from the previous block in face space (e.g., the current block was F1 and the previous block was F6), this regressor would be assigned a value of 0.2, corresponding to the potentially small contribution of adaptation. A value of zero was assigned when the current block was preceded by a fixation block. All results were largely the same whether or not this regressor was included.

For the face attention task (Experiment 2), a single GLM was computed for all runs, with separate regressors for each of the three conditions (attend F1, attend F6, and fixation) and for the response period. In two sessions, the response period was not in a separate volume from the stimulus block and was therefore not included as a separate regressor.

Region of interest selection. For each session, face areas were defined based on a contrast of faces versus scenes from the face localizer task ($p < 0.001$, uncorrected; this somewhat liberal threshold was used since a second voxel selection step was included in computing tuning curves). Voxels in temporal and occipital cortex that responded significantly more to faces than to scenes were included in a single “all face areas” region for primary analyses. This region included clusters of voxels within the OFA, FFA, and fSTS, as well as several small clusters or single voxels, mainly in the anterior temporal lobe, that were combined and labeled as anterior temporal areas (fAT; Fig. 2). Collectively, these areas correspond to regions classically identified as part of the face processing network (Haxby et al., 2001).

Voxel tuning curve procedure. We adapted a method described by Serences et al. (2009) to classify the face preference and construct face tuning curves for individual voxels. For each face-selective (preferring faces to scenes) voxel identified with the face localizer, separate β values were

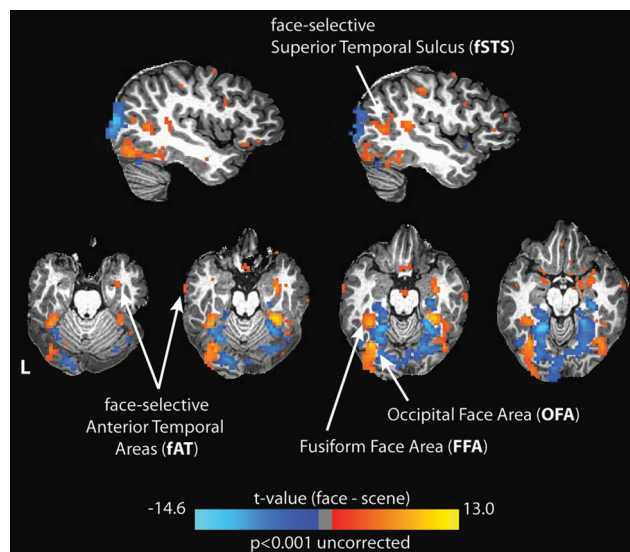


Figure 2. Identification of face areas. Face-selective regions in temporal and occipital cortex were defined in each session using a localizer (faces > scenes; displayed in warm colors). An example from one subject is shown here.

extracted for each of the six faces in the morph continuum for each face mapping run, using the contrast of each face condition (F1 through F6 blocks) versus fixation. These β values were normalized within each voxel by subtracting the voxel’s mean response across all six face conditions. This was done to remove differences across voxels in their overall responsiveness, thereby facilitating measurement of differences in the pattern of responses to different faces within a voxel. We then used all but one face mapping run to classify each voxel as preferring one of the six face morphs, based on the face that elicited the highest average β value across these classification runs. The left-out run that was not used in classification of the preferred face for each voxel was then used to measure responses to each of the six face types, creating a tuning curve for each voxel. This leave-one-out procedure was repeated with all possible combinations of classification and analysis runs for cross-validation.

A further selection step was applied during this procedure: only those voxels that had the same preferred face in more than half of the classification runs were considered to be “reliable” voxels and included in further analysis (on average, 22 voxels per session for all face areas combined). Due to the small size of the face-selective anterior temporal area, this reliability criterion was relaxed to a threshold of $>20\%$ agreement in preferred face across classification runs, thereby resulting in similar numbers of reliable voxels for posterior and anterior face areas.

Tuning curves were averaged across all voxels with the same face preference (F1 through F6) and across all runs of the cross-validation procedure. Because the number of reliable voxels varied widely across individual sessions, group analyses were weighted by the number of voxels that each session contributed to the analysis using the procedure described by Bland and Kerry (1998). In addition, to examine whether some areas had tuning that was more categorical in nature rather than specific for individual faces along the morph continuum, the same procedure described above for classifying voxels based on individual face preference was used to classify voxels as preferring either the first half (F1–F3, the “male” half) or the second half (F4–F6, the “female” half) of the face morph continuum. For the categorical classification, β values were averaged across the three conditions within each half of the face continuum (i.e., averaged responses to either F1 through F3 or F4 through F6) for each voxel.

Statistical analyses of tuning curves. To assess how reliably our tuning curve procedure classified the preferred face type for individual voxels, weighted β values from each session were entered into a repeated-measures ANOVA with two factors: voxel type (F1–F6) and condition (preferred vs all nonpreferred face types, e.g., for F1 voxels, F1 condition responses vs the average of F2–F6 condition responses). This and all

other ANOVAs were conducted using the ezANOVA package of R (R Development Core Team, 2011). The p values for the main effect of condition in the ANOVA are the results of one-tailed tests, reflecting the a priori directional hypothesis that preferred conditions should elicit higher responses than nonpreferred conditions. These results were further examined with *post hoc* one-tailed t tests comparing responses of preferred to nonpreferred conditions for each of the classified voxel types. These same methods and statistical procedures were also applied to secondary analyses examining the tuning separately in posterior and anterior face areas and in motor cortex and early visual cortex control regions for tuning to all six individual faces as well as for tuning to face categories (F1–F3 vs F4–F6).

The weighted β values from each session were also used to calculate the intraclass correlation coefficient (ICC) between the two fMRI sessions from the same participant. A separate ICC value was computed for each of the six voxel types, and these ICC values were averaged. Mean ICC values were -0.07 , indicating a low correspondence at the individual voxel level between fMRI sessions for the same subject. We therefore analyzed all data separately for each session ($N = 20$). It is unsurprising that ICC values would be low, since face preference was determined separately for each voxel, and the exact population of neurons in a voxel is likely to be strongly affected by the specific head position and slice prescription in each session.

In addition, to probe whether tuning curves changed systematically along the morph continuum, we conducted a session-specific linear regression analysis on the difference between F1 and F6 tuning curves (F1 minus F6). These correlation (r) values were normalized by Fisher transformation, and a one-tailed paired t test was then used to test for significant differences from zero (reflecting the a priori hypothesis that a negative relationship should exist; i.e., the difference between F1 and F6 responses should be more positive when F1 is presented and more negative when F6 is presented). Note that here, unlike in the tuning curve analysis, the values were not weighted by the number of reliable voxels in each session, as multiple voxel types were combined for this analysis.

Additional analyses quantified variability in the number of voxels representing different faces. The percentages of reliable voxels classified as preferring each face were entered into a one-factor repeated-measures ANOVA. Here, unlike the other ANOVAs, more than two groups were compared, so statistical results were corrected for sphericity using the Huynh–Feldt method in the ezANOVA package of R (R Development Core Team, 2011).

Analysis of attentional modulation. Beta values from the face attention task were computed for each reliable voxel and examined according to each voxel's preferred face type (defined by the classification procedure on the independent face mapping data set from Experiment 1, but using all mapping runs, rather than one run left out). The differences in β values between attend-F1 (compared with fixation) and attend-F6 (compared with fixation) contrasts were then calculated for each of the six voxel types. We tested for a U-shaped response pattern with two-tailed paired t tests that compared responses in voxels preferring one end of the morph continuum to responses in voxels preferring intermediate faces (F1 vs F2–F5 and F6 vs F2–F5).

To directly compare the two attention conditions, β values from the attend F6 condition were subtracted from those from the attend F1 condition after subtracting the average response to each condition (the attend F6 condition elicited overall higher responses, perhaps contributing to the slightly higher morph detection rate observed in this condition). We tested for attentional modulation of responses with paired two-tailed t tests on the normalized β values within each voxel type. These p values were corrected for multiple comparisons using the false discovery rate (FDR) procedure, implemented in the p.adjust program of R. An analogous procedure was used to analyze attention affects in voxels classified categorically (F1–F3 vs F4–F6).

Control analyses. To determine whether voxel-based face tuning curves were only present in face-selective (i.e., preferring faces to scenes) regions, we conducted tuning curve and attentional modulation analyses in a control area comprised of bilateral motor cortical regions. This area was selected from the FSL Harvard Oxford Probability atlas [thresholded at 25% probability (Desikan et al., 2006)] and reverse normalized into each

Table 1. Number and size of face areas examined

	Number of sessions (N)	Voxel number (mean \pm SEM)	Voxel count range
All face areas	20	2000 \pm 341	113–5541
Posterior face areas (OFA, FFA, fSTS)	18	1636 \pm 318	77–5400
Anterior face areas (fAT)	17	267 \pm 95	13–1507
Motor cortex (control)	20	7218 \pm 457	2415–8199
Early visual cortex (control)	20	2022 \pm 154	636–3325

The number of sessions for which each face-selective region could be identified and the mean size of each region are listed. Regions examined included posterior face areas (OFA, FFA, and fSTS), anterior face areas (small clusters in the anterior temporal lobe, fAT), and two control sites (motor cortex and early visual cortex). In most cases, regions were bilateral.

individual subject's native space. The control region was then masked to remove nonbrain voxels (leaving an average of 7218 voxels, SEM = 457; Table 1). All analyses described above (including selection of reliable voxels, classification, tuning, and attention analyses) were then conducted for this control region. This constitutes a strong test of possible bias in our procedures and of the specificity of our findings to visual regions, since using a large area with many voxels increases the probability of finding some voxels with face tuning.

We also conducted a second control analysis using an area centered over the calcarine sulcus and covering early visual cortex. This region was selected from the FSL Harvard Oxford Probability atlas [intracalcarine, thresholded at 25% probability (Desikan et al., 2006)]. In addition, because this region did not extend to the occipital pole (and might, therefore, have excluded foveal visual field representations within the stimulus), we manually drew an extension for this region to the occipital pole. We then reverse normalized this region into each individual subject's native space and masked it to remove any nonbrain voxels (leaving an average of 2022 voxels per session, SEM = 154; Table 1). Note that this region is similar in size to the combined posterior and anterior face-selective area used for our main experimental analyses. As before, all previously described analyses were conducted for this early visual control region, constituting a strong test of the specificity of our findings to higher-level visual regions.

Results

Subjects participated in two experiments within the same fMRI session (Fig. 1), and each fMRI session was analyzed separately (see Materials and Methods). The aim of Experiment 1 was to generate voxel-based tuning curves for a set of six morphed faces (F1–F6), and the aim of Experiment 2 was to measure the effects of feature-based attention to one of the previously presented faces (F1 or F6) on these voxel-based tuning curves.

Identification of face selective regions

In each session, clusters of voxels in temporal and occipital cortex that responded more to faces than scenes during the localizer task were selected for further analysis (Table 1). These clusters included regions classically identified as part of a face processing network (Haxby et al., 2000), including the FFA, OFA, fSTS, and regions in the anterior temporal lobe (fAT; see Fig. 2 for an example from one participant and Table 1 for the size and number of face areas that were identified across sessions).

Experiment 1: face tuning

Behavior

In Experiment 1, blocks of a single face morph (F1, F2, F3, F4, F5, or F6) or fixation were presented to participants while they counted the number of contrast decrements within each face morph block (Fig. 1). This encouraged them to maintain spatial attention on the face without attending to specific facial identities or features. The amount of contrast decrement was adjusted for each run to maintain accuracy at $\sim 75\%$ correct across runs

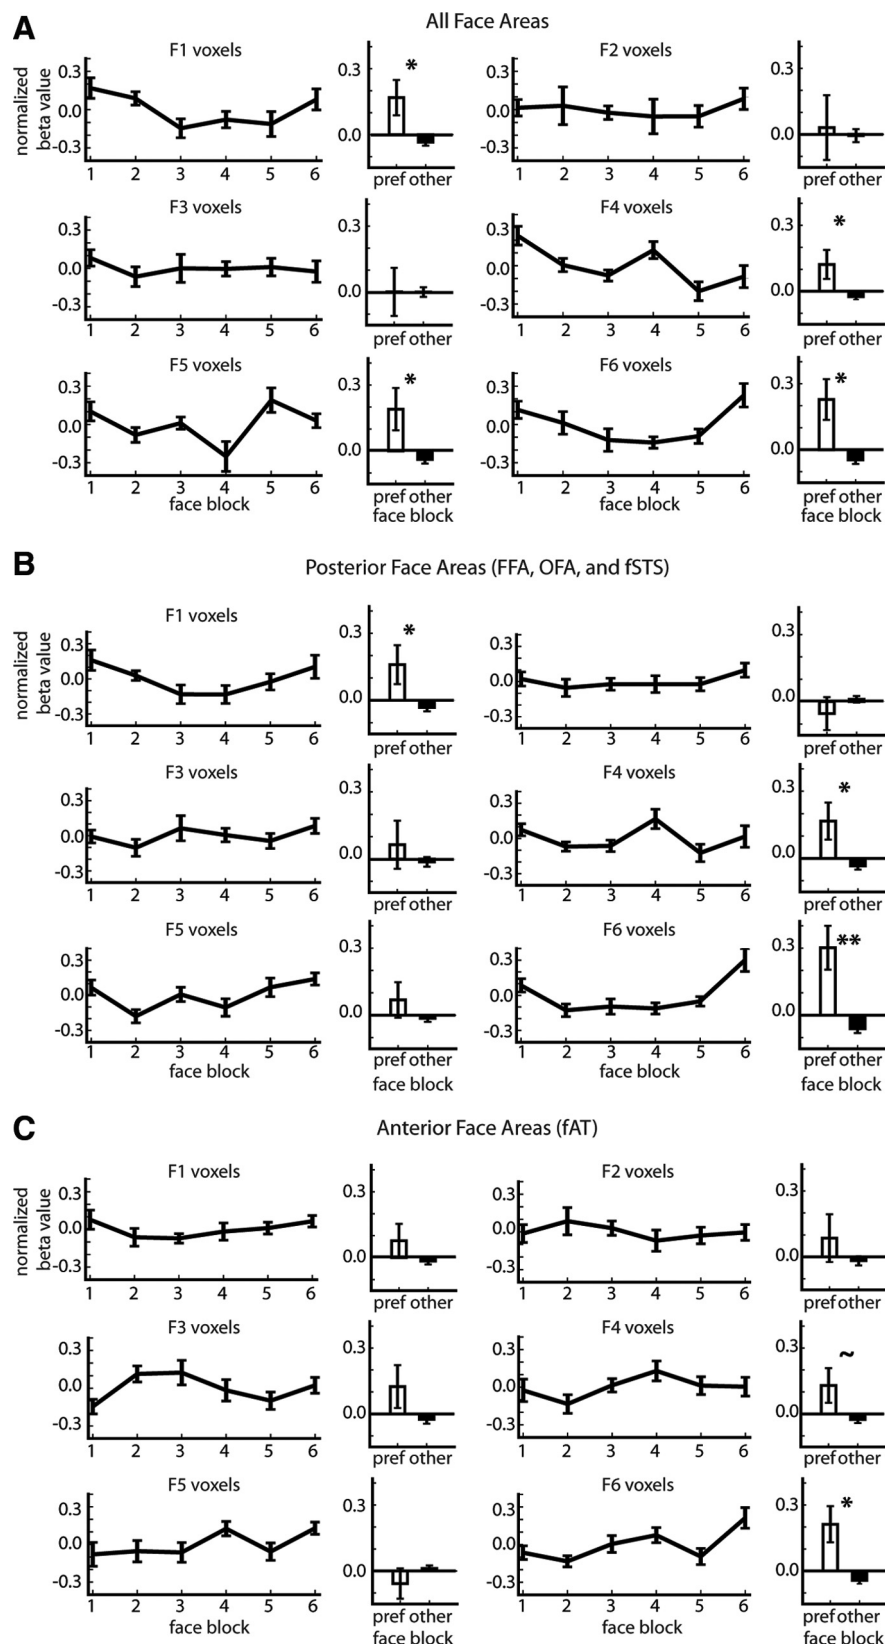


Figure 3. Voxel tuning curves (Experiment 1). **A**, Voxel-based tuning curves are shown for all reliable face-selective voxels in occipital and temporal cortex. Tuning curves for voxels preferring each face morph (F1–F6) are displayed in separate subplots across the six face conditions (error bars = SEM across fMRI sessions for this and all subsequent figures). The bar graphs to the right of each tuning curve plot the response to the preferred face compared with the average response to nonpreferred faces for that voxel type (e.g., for F1 voxels, the response to the F1 condition vs the mean response to F2–F6 conditions). Data are also displayed for voxels in only posterior (**B**) and anterior (**C**) face areas. * $p < 0.05$; ** $p < 0.01$; ~ $p < 0.10$.

(chance performance = 25%). Mean accuracy on the task was 73.4% (SD = 6.6%), and the mean contrast decrement required to maintain this accuracy was 0.122 (SD = 0.04).

Voxel-based tuning across all face-selective areas

We constructed voxel-based tuning curves (Fig. 3) of responses to each face morph in the continuum by using a cross-validation procedure described by Serences et al. (2009) (see Materials and Methods). This method classifies the preferences of individual voxels directly from the amplitude of the hemodynamic responses to presentation of face stimuli and is distinct from methods that classify voxels based on their attenuation to repetition, as in fMR-A, or on the weights they contribute to a model, as in MVPA studies. Critically, this technique allows us to determine response profiles of single voxels (rather than across a population, as with MVPA) that can then be compared across experimental conditions.

We assessed the reliability of our voxel-based tuning curves using a two-factor random effects ANOVA, with factors of voxel type (F1–F6) and condition (preferred or nonpreferred; e.g., for a voxel classified as F1, the preferred condition was responses in the F1 condition, and the nonpreferred condition was the average of responses in the F2–F6 conditions). As a group, tuning curves demonstrated selectivity for the preferred face compared with nonpreferred faces ($F_{(1,19)} = 4.20$, one-tailed $p = 0.027$). There was no significant main effect of voxel type or interaction between voxel type and condition. *Post hoc t* tests revealed that four of the six tuning curves had significant selectivity for the preferred face (voxels preferring F1: $p = 0.024$; F4: $p = 0.041$; F5: $p = 0.031$; F6: $p = 0.012$; Fig. 3A).

Tuning in posterior versus anterior face areas

Next, tuning curves were examined separately for voxels in more posterior (FFA, OFA, fSTS) (Fig. 3B) and anterior (fAT) (Fig. 3C) face areas. As a group, tuning curves in posterior areas demonstrated significant selectivity for the preferred face ($F_{(1,17)} = 3.87$, one-tailed $p = 0.033$). In anterior areas, tuning curves demonstrated only marginally significant selectivity for the preferred face ($F_{(1,15)} = 2.86$, one-tailed $p = 0.056$). *Post hoc t* tests showed that in posterior areas, three of six tuning curves showed significant selectivity (F1: $p = 0.042$; F4: $p = 0.030$; F6: $p = 0.003$), while only one tuning curve showed significant selectivity in the ante-

rior areas (F6: $p = 0.010$; marginally significant in F5: $p = 0.060$). There was no significant main effect of voxel type or interaction between condition and voxel type in either set of areas.

Additional tuning analyses

To further assess whether tuning differed systematically across the morph continuum, we computed the difference between the F1 and F6 within-voxel tuning curves (F1 minus F6) and conducted a linear regression analysis on data from each session. If tuning curves change smoothly along the morph continuum, there should be a systematic progression from stronger responses in F1 voxels for faces similar to F1 to stronger responses in F6 voxels for faces similar to F6, resulting in a negative slope in the difference between the F1 and F6 tuning curves. In posterior areas, this negative linear relationship between the location of the presented face within the morph continuum and the F1/F6 response difference across sessions was observed ($r = -0.22$, $p = 0.035$), but this was not the case for anterior areas ($r = -0.03$; Fig. 4A). This indicates that in posterior areas, response amplitude varied systematically as a function of the six morphed faces, with higher responses in F1 voxels for faces closer to F1 within the face continuum and higher responses in F6 voxels for faces closer to F6.

It is possible that in some brain areas, faces are represented in a more categorical fashion (e.g., F1 through F3 vs F4 through F6) rather than being represented as individual exemplars of faces along a morph continuum. This sort of division might allow for representations based on different categories within this set of face morphs, such as male versus female or old versus young. We therefore used the same cross-validation procedure to recompute the voxel tuning curves for only two classes of voxels: male voxels (preferring F1 through F3 more than F4 through F6) and female voxels (preferring F4 through F6 more than F1 through F3; note that use of the terms male and female here is somewhat arbitrary since the two halves differed along multiple categorical dimensions). This binary classification revealed significant tuning in anterior ($F_{(1,15)} = 5.77$, one-tailed $p = 0.0149$) but not in posterior ($F_{(1,17)} = 1.10$) face areas (Fig. 4B). *Post hoc t* tests demonstrated that in anterior areas, both male and female voxels showed significantly higher responses to their preferred face class relative to their nonpreferred face class (first half: $p = 0.040$; second half: $p = 0.012$). This indicates that, although anterior face areas only show marginally significant tuning to individual face morphs, they exhibit significant tuning when comparing responses to the two halves of the face continuum. That is, while posterior areas primarily show tuning for individual faces along the entire morph continuum, anterior areas mainly show tuning for face categories within this continuum.

Experiment 2: face attention

In addition to the face mapping task, participants completed a feature-based attention task in the same session (Fig. 1B). This allowed for analysis of the effects of attention on voxel-based

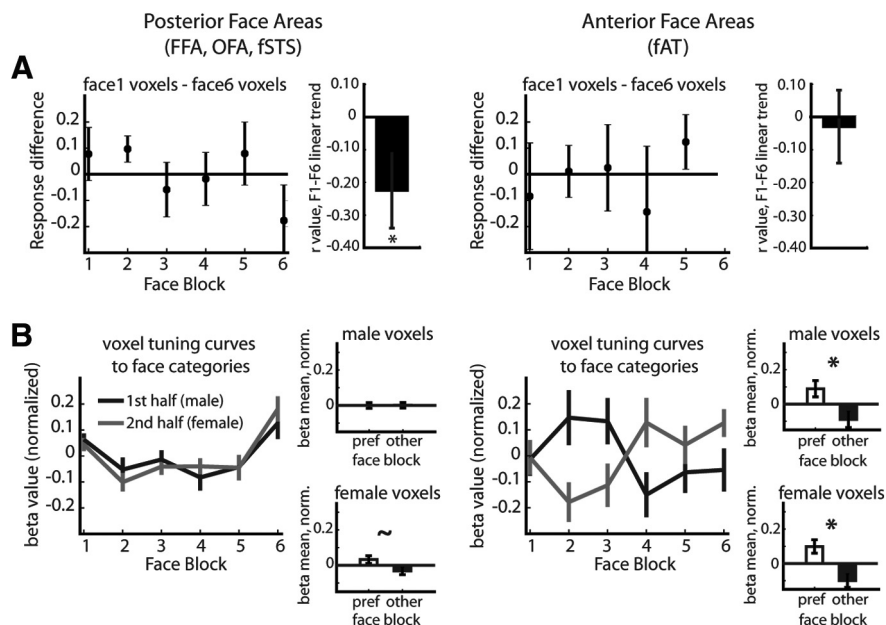


Figure 4. Posterior and anterior face areas have different tuning properties (Experiment 1). **A**, A regression analysis was conducted on the difference between tuning curves for F1 and F6 voxels in individual fMRI sessions. A significant linear relationship was present in posterior but not anterior face areas. **B**, Tuning curve plots show responses for voxels classified into two face categories (black, F1 through F3; gray, F4 through F6). Bar graphs compare responses to preferred and nonpreferred conditions (e.g., for male voxels, preferred = F1, F2, and F3 blocks; nonpreferred = F4, F5, and F6 blocks). Responses to the preferred half of the face morph continuum were significantly higher than responses to the nonpreferred half in anterior (both halves $p < 0.05$) but not in posterior face areas.

tuning for faces. Participants viewed superimposed images of F1 and F6 and were instructed to attend to just one of the two faces in each block and to count the number of morphs in the attended face toward a third face.

Behavior

The amount of morph toward the third face in each block was modified on each run for each attention condition to maintain morph detection performance at $\sim 75\%$ correct across runs (chance performance = 25%). Mean accuracy was 71.3% (SD = 5.5%) for attending to F1 and 74.8% (SD = 6.4%) for attending to F6. On average, a 34.7% (SD = 7.4%) morph for F1 and 35.0% (SD = 10.4%) morph for F6 were required to maintain this performance level.

Attention effects across all face areas

We used the voxel classifications from Experiment 1 to measure attentional modulation of reliable face-selective voxels to a composite display consisting of only the two faces at either end of the face morph continuum (F1 and F6), superimposed and tilted 45 degrees in opposite directions. Therefore, if the voxel classification procedure from Experiment 1 was accurate, responses to the composite display (relative to fixation) should be elevated in voxels preferring F1 and F6, compared with voxels preferring intermediate faces. Indeed, responses to the composite F1/F6 image during the attention task exhibited a U-shaped profile across voxel types in both attend-F1 and attend-F6 conditions, with highest responses in voxels classified as preferring F1 and F6 in Experiment 1, and lower responses in F2–F5 voxels (Fig. 5A; F1 vs F2–F5: $t_{(19)} = 2.84$, $p = 0.011$; F6 vs F2–F5: $t_{(19)} = 3.84$, $p = 0.001$). Therefore, classification of voxels from Experiment 1 correctly predicted which voxels would exhibit the highest response in an independent data set from Experiment 2. This corroborates the findings from Experiment 1, providing further evidence that

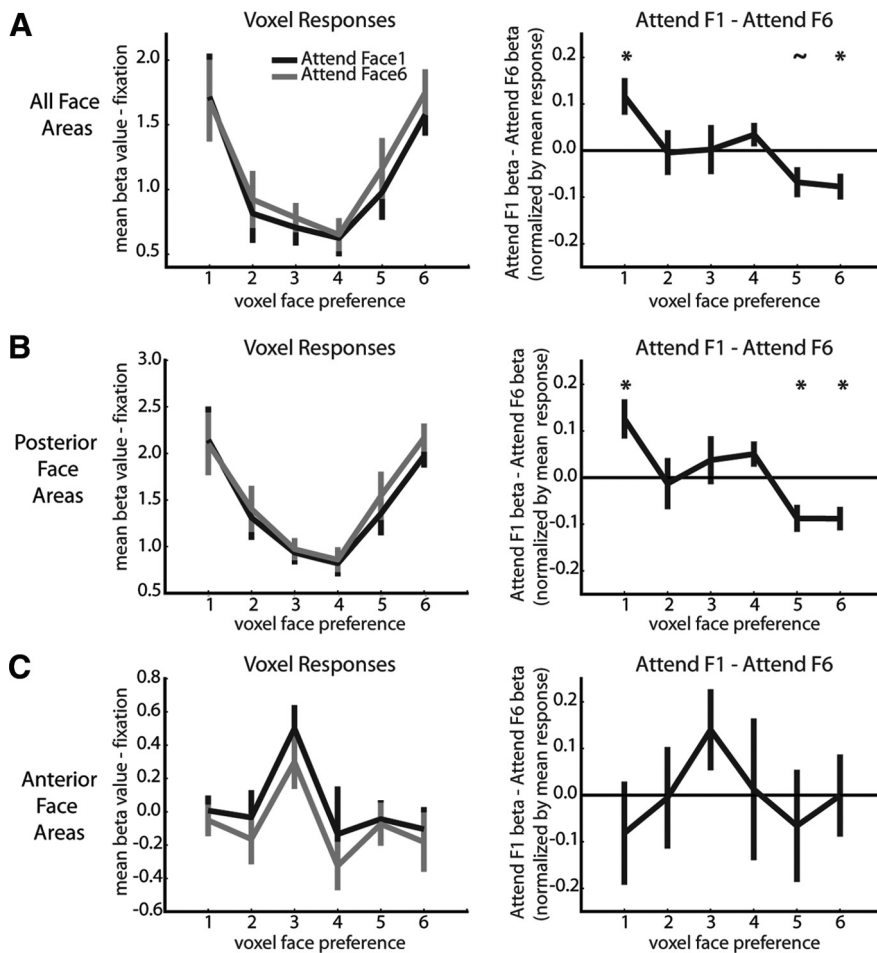


Figure 5. Responses during attention task (Experiment 2) across voxel types. **A**, Left, Beta values for each attention condition (compared with fixation) from voxels across all face areas are plotted according to the face preference those voxels displayed in Experiment 1. **A**, Right, The difference between attention conditions (normalized by subtracting the mean response to each condition across all voxel types) is shown for voxels across all face areas. The same information is plotted for voxels in only posterior (**B**) and anterior (**C**) face areas. Note that in this figure (unlike Figs. 3, 4), responses are plotted across voxel types on the x-axis rather than across stimulus conditions (within voxel types). * $p < 0.05$; ~ $p < 0.10$; FDR-corrected.

our voxel classification procedure correctly identified voxel preferences for individual faces. This concordance of the face mapping and attention data is notable, since it occurred despite differences between the two experiments in task, size of stimulus (in degrees of visual angle), and stimulus configuration.

To assess the effects of feature-based attention on responses to the pair of superimposed faces, we examined the difference between conditions (attend F1 minus attend F6, following subtraction of the average response to both attention conditions in Experiment 2, across all voxels that were classified as reliable in Experiment 1). Attending to F1, relative to attending to F6, selectively enhanced responses in F1 voxels (Fig. 5A; $t_{(19)} = 2.97$, $p[\text{FDR}] = 0.032$). Similarly, attending to F6, relative to attending to F1, selectively enhanced responses in F6 voxels ($t_{(19)} = -2.84$, $p[\text{FDR}] = 0.032$) and marginally enhanced responses in F5 voxels ($t_{(19)} = -2.13$, $p[\text{FDR}] = 0.092$).

Attention effects in posterior versus anterior face areas

Posterior areas showed a similar pattern of attention effects to those seen for all face areas combined. Responses were U-shaped across voxel types, again suggesting that classification based on face mapping data correctly identified the six classes of voxels (Fig. 5B; F1 vs F2–F5: $t_{(17)} = 2.75$, $p = 0.014$; F6 vs F2–F5: $t_{(17)} =$

4.30, $p = 0.0005$). Moreover, in posterior areas, attending to F1 selectively enhanced responses in F1 voxels ($t_{(17)} = 3.01$, $p[\text{FDR}] = 0.017$), and attending to F6 selectively enhanced responses in F5 and F6 voxels (F5: $t_{(17)} = -3.06$, $p[\text{FDR}] = 0.017$; F6: $t_{(17)} = -3.52$, $p[\text{FDR}] = 0.016$) (Fig. 5B).

In anterior face areas, the response profile was not U-shaped (Fig. 5C; F1 vs F2–F5: $t_{(15)} = -0.351$, F6 vs F2–F5: $t_{(15)} = -0.429$), and there were no significant differences between attention conditions (Fig. 5C; F1: $t_{(15)} = -0.67$; F6: $t_{(15)} = 0.11$, all $p[\text{FDR}] > 0.17$). In addition, anterior face areas exhibited no significant attention effects even when responses were categorically grouped by voxels preferring each face class (F1 through F3 and F4 through F6) rather than by voxels preferring individual face morphs (F1 through F3: $t_{(15)} = -1.75$, F4 through F6: $t_{(15)} = 1.68$). Note that numerically, these values in anterior face areas are in the opposite direction from that expected if attention to a feature enhanced responses in neurons selective for that feature.

Experiments 1 and 2: control area tuning and attention

To assess whether there was any bias in our methods for constructing voxel-based tuning for faces or for measuring attentional modulation of this tuning, we conducted all analyses on an anatomically defined bilateral motor cortical region. Despite the large size of this region, which should increase the probability of finding a small subset of voxels with tuning, voxel tuning across the six faces was poor (Fig. 6A). The main effect of condition (preferred vs nonpreferred) was not significant ($F_{(1,19)} = 0.448$), and there was no detectable main effect of voxel type or interaction between voxel type and condition. Only one of the six tuning curves showed significantly higher responses for preferred, compared with nonpreferred, faces in *post hoc* t tests (F1: $t_{(19)} = 3.0$, $p = 0.004$). Furthermore, there was no significant attentional modulation in any voxel type (Fig. 6B, bottom; F1: $t_{(19)} = 0.47$, F6: $t_{(19)} = 2.19$, note that the F6 value in motor cortex is in the opposite direction from that expected if attention to a feature enhanced responses in neurons selective for that feature).

In addition, we conducted a separate control analysis in an anatomically defined early visual cortex region to assess the specificity of our results to higher-level visual regions. We observed strong voxel tuning to the six faces in Experiment 1 in this region (Fig. 6C), and the main effect of condition was highly significant ($F_{(1,19)} = 27.52$, one-tailed $p = 0.00002$). In addition, four of the six individual tuning curves had significantly higher responses to preferred, compared with nonpreferred, faces, based on *post hoc* t tests (voxels preferring F1: $p = 0.0001$; F2: $p = 0.006$; F5: $p = 0.019$; F6: $p = 0.012$; F3: $p = 0.052$). However, the U-shaped response profile was not present in Experiment 2 (Fig. 6D, top, F1 vs F2–F5: $t_{(19)} = 4.73$, $p = 0.0001$; but F6 vs F2–F5: $t_{(19)} = 0.32$), suggesting limited invariance to the dif-

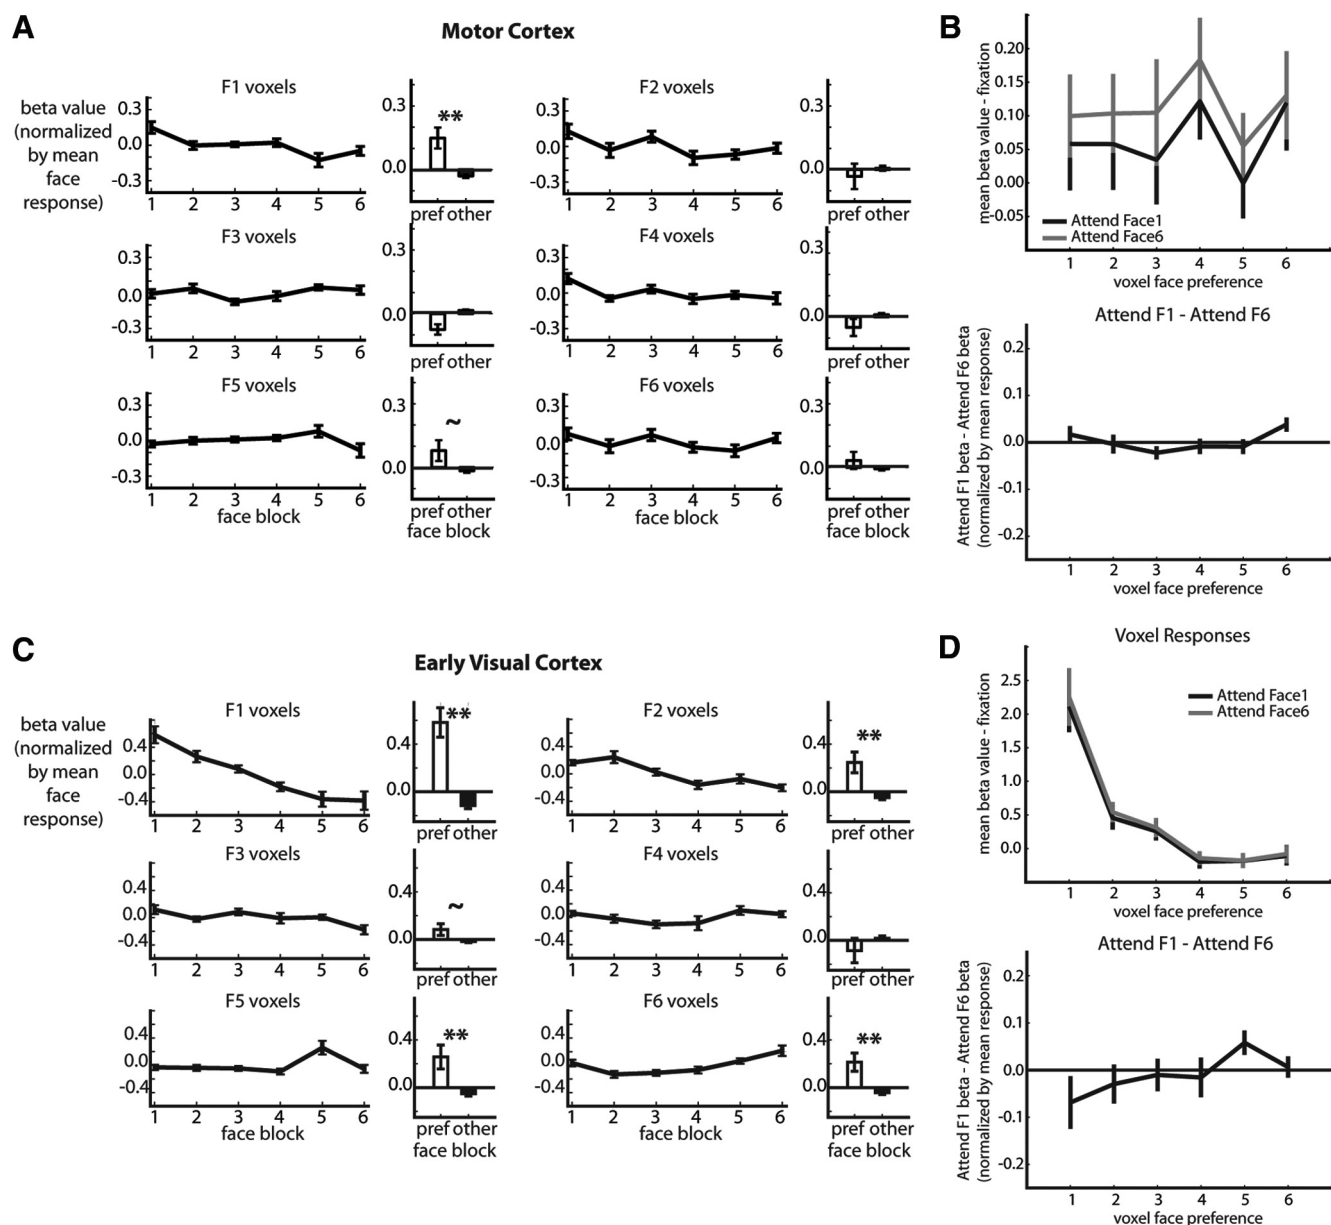


Figure 6. Tuning and attention results in control bilateral motor and early visual cortical regions. Group voxel-based tuning curves for bilateral motor (**A**) and early visual (**C**) regions, plotted using conventions from Figure 3 (** $p < 0.01$; $\sim p < 0.10$; note that the scale is larger in subplots for **C** to accommodate the greater range of β values). Responses in the two attention conditions for motor (**B**) and early visual (**D**) areas are displayed on the top subplot using the same conventions as Figure 5. The difference between attention conditions, normalized by the average response to each condition, is displayed in the bottom subplot for each area, as in Figure 5 (all n.s.).

ferences in stimulus properties between Experiments 1 and 2, including contrast and orientation of the faces. Furthermore, there was no significant attentional modulation in any voxel type (Fig. 6D, bottom, F1: $t_{(19)} = -1.23$, F6: $t_{(19)} = 0.28$).

Discussion

We used a method previously used to construct voxel-based tuning curves for low-level features (Serences et al., 2009) to identify individual voxels that were selective for faces along a morph continuum (F1–F6). Different cortical areas showed distinct patterns of tuning: posterior face areas (FFA, OFA, fSTS) were primarily selective for individual face morphs, while more anterior face areas mainly exhibited categorical tuning for faces (F1–F3 vs F4–F6, e.g., male vs female halves). Data from a separate task in which subjects attended one of a pair of superimposed faces (F1 and F6) validated these tuning

results, showing that F1 and F6 voxels exhibited significantly larger responses to a stimulus containing both faces, compared with voxels preferring intermediate face types. Critically, directing attention to a particular face selectively enhanced responses in voxels previously defined as preferring that face. Our results demonstrate that fMRI can be used to classify individual face preferences in single voxels and provide direct evidence that the basic principles of selectivity of attentional effects that have previously been described for lower level features apply equally well to attentional modification of the representations of complex stimuli.

Relationship to previous studies of neural face representations

Our findings complement previous studies that examined neural selectivity for faces across different spatial scales. Single neuron

responses can be selective for face stimuli (for review, see Gross, 2011), face feature dimensions and feature combinations (Freiwald et al., 2009), face morph vectors (Leopold et al., 2006), and even individual identities (Quiroga et al., 2005). The selectivity for individual face stimuli that we observed may depend on a combination of signals that are tuned for multiple distinct elements (individual faces, face features, and face morphs).

In addition, fMR-A has been used to show that responses in face-selective areas are sensitive to differences between individual faces (Winston et al., 2004; Pourtois et al., 2005; Cohen Kadosh et al., 2010), face parts (Harris and Aguirre, 2008; Andrews et al., 2010), and face morphs (Loffler et al., 2005; Rotshtein et al., 2005; Gilaie-Dotan and Malach, 2007). However, fMR-A measures may not accurately reflect neuronal selectivity [due to a number of factors, including possible feedback/feedforward influences (Kriegeskorte et al., 2007; Bartels et al., 2008; Mur et al., 2010) and a lack of correspondence between neuronal response selectivity and neuronal adaptation selectivity (Sawamura et al., 2006)]. In contrast, our study used a method that is not susceptible to these issues and directly demonstrates voxel-level selectivity for individual faces.

It has also been suggested that multivoxel pattern classification methods can be used to reveal a distributed representation of high-level features (Haxby et al., 2001). These methods have recently identified fMRI responses differentiating individual faces (Kriegeskorte et al., 2007; Natu et al., 2010; Nestor et al., 2011). Our results similarly demonstrate that voxels from multiple areas contain signals that are selective for individual faces. However, we found that only a small proportion of voxels within these areas were both reliably selective for the presented face morphs and modified by attention. This may not be surprising, as the signals that drive classification presumably result from small differences in the relative proportions of neurons tuned to different faces within each voxel, and we only tested a small number of faces in our study. Our results extend previous MVPA studies by demonstrating that single voxels are tuned for facial identity.

As noted above, our results agree with previous electrophysiological and fMRI studies investigating neural tuning for faces. However, since we only tested a small set of individual face stimuli, it is difficult to determine the extent to which the tuning properties described here generalize to a larger set of stimuli. Future studies with larger stimulus sets can add to these findings by more precisely determining the stimulus parameters to which individual voxels are tuned.

Tuning curve properties in different face areas

Our results demonstrate that face areas vary in the relative levels of continuous featural codes versus categorical codes used to represent face stimuli. Specifically, posterior face areas showed selectivity for individual face morphs. Furthermore, the difference in tuning between F1 and F6 voxels indicated a significant monotonic relationship, suggesting that tuning varied systematically along the morph continuum. In contrast, anterior face areas showed selectivity for face categories (i.e., F1 through F3 vs F4 through F6, corresponding to male vs female halves of the continuum).

Our results support existing models of differences in representations across face areas. Specifically, anterior temporal face areas have been proposed to connect physical face information with more semantic, biographical detail, with posterior regions (OFA, FFA, and STS) being more involved in the processing of face features, face individuation, and gaze and expression information, respectively (Haxby et al., 2000). Neurophysiological

studies also suggest differences in face coding along the posterior–anterior axis of the temporal lobe, with neurons coding for face features in more posterior regions (Freiwald et al., 2009) and for face morph vectors in more anterior regions (Leopold et al., 2006). Work in patients with prosopagnosia is also consistent with this distinction: patients with damage to posterior temporal or occipital cortex may be impaired in perceiving facial identity based on visual information, whereas patients with damage to the anterior temporal lobe are more likely to be unable to recognize an individual regardless of the source of information (Damasio et al., 1990).

One potential interpretation of our findings is that posterior areas show selectivity for individual faces, whereas selectivity in anterior face areas is related to broad classes of faces (e.g., male vs female, old vs young). However, previous psychophysical studies (Beale and Keil, 1995) and fMRI data from some face-selective regions (Rotshtein et al., 2005) suggest that individual face morphs along a continuum are not necessarily each represented as distinct identities, but rather as a set of two identities with a categorical boundary at some point along the morph continuum. Therefore, our findings in posterior areas could also be interpreted as these regions showing sensitivity to face features (rather than identities) that differed between individual morphs, whereas our findings in anterior areas may reflect sensitivity to distinct facial identities. This interpretation of our findings adds support to past literature by providing population-tuning evidence for a more feature-based code (selectivity for individual morphs) in posterior cortex and a more categorical code (selectivity for halves of the morph continuum) in anterior cortex. Although it is beyond the scope of this paper, future work could examine whether this category/exemplar distinction is represented as a continuous gradient along the anterior–posterior axis of the inferior temporal lobe or if there is a sharp boundary in this region that demarcates distinct coding mechanisms in different face areas.

Feature-based attention modulates voxel-level face representations

We have found that attention can selectively modify responses in voxels tuned for individual exemplars within a category. Previous studies have suggested that feature-based attention increases responses of neurons to attended features in lower-level visual regions such as V4 (Motter, 1994; McAdams and Maunsell, 2000) and MT (Treue and Martínez Trujillo, 1999; Martínez-Trujillo and Treue, 2004). Our finding that attention selectively enhances responses in voxels representing the attended face is consistent with an increase in population gain. However, it has also been suggested that attention can shift and/or reduce the size of receptive fields of individual visual cortical neurons (Connor et al., 1996; Womelsdorf et al., 2006; David et al., 2008) and sharpen the tuning of populations of neurons [through feature-similarity gain in individual neurons (Martínez-Trujillo and Treue, 2004)]. The strongly selective nature of the attentional enhancement that we observed is also consistent with modifications in population tuning of responses. We found that attention amplified responses only in voxels selective for the attended face (or sometimes in voxels preferring the face adjacent to the attended face along the morph continuum: F5 when attending to F6). This selectivity of attention effects occurred despite broad individual tuning curves in different voxel types.

In contrast to the extensive work on attentional modulation of tuning for low-level features, considerably less is known about attention's effects on tuning for complex stimuli. Previous stud-

ies have demonstrated that attending to items within a category can enhance signals across entire category-selective brain areas (Corbetta et al., 1991; Wojciulik et al., 1998; Gazzaley et al., 2005), shift MVPA signals toward the attended category (Reddy et al., 2009; Chen et al., 2012), and cause sharper and larger release from adaptation with fMR-A in these areas (Murray and Wojciulik, 2004). By showing that attention can operate selectively on individual exemplars within a complex category, we establish a connection between the effects of attention on complex objects and on low-level visual representations (Tootell et al., 1998; Brefczynski and DeYoe, 1999; Somers et al., 1999; Serences et al., 2009; Saproo and Serences, 2010), thereby providing evidence that the effects of attention on feature representations are precisely tuned across the visual hierarchy.

Attentional modulation of different face selective visual areas

We observed feature-based attentional modulation across all face areas when they were grouped together and in posterior but not anterior areas when examined separately. This dissociation may be because the attention task we used emphasized subtle shifts in facial features (small morphs toward a face not on the morph continuum). We found that posterior visual areas were sensitive to small changes in stimulus properties (i.e., differences between individual faces along the morph continuum), whereas anterior areas responded to larger categorical changes. Therefore, attention may have selectively acted on those brain areas that were most relevant for performing the task. Another possibility is that the smaller number of face-selective voxels within the anterior temporal lobe (Table 1) resulted in decreased reliability of population responses of this area, thereby limiting our ability to detect attentional modulation. This interpretation is supported by the absence of a U-shaped profile of responses in the face attention experiment (Fig. 5C).

Specificity of tuning and attention effects to face-selective areas

To determine whether our results were unbiased and selective to face-selective areas of cortex, we performed a control set of analyses in bilateral motor and early visual cortical regions. Despite the large size of the motor cortex region, which could have favored detection of a subset of voxels with appropriate tuning properties, we found no evidence that voxels responded selectively to individual faces within this region. Furthermore, we found no evidence for attentional modulation of responses in the motor cortex control region. These results show that our method for identifying voxels tuned for individual faces is unbiased and that the tuning and attentional modulation we report are specific to visual areas.

In addition, we used an early visual cortex control region to further investigate the specificity of our findings to higher-level visual cortex. In this region, we found strong evidence for selective tuning to individual faces in Experiment 1, an expected result given that our stimuli differ from one another in a number of low-level features. In fact, there is evidence from previous studies for very selective tuning to natural images in early visual cortex (Kay et al., 2008). However, the lack of a U-shaped response profile in Experiment 2 suggests that this tuning is not invariant to small changes in the stimulus configuration (e.g., rotating the faces 45 degrees, changing their contrast). This is precisely what one would expect in an area with small receptive fields tuned to low-level features. Finally, and most importantly, the attention effects we observe in higher-level face-areas were absent in the early visual control area, suggesting that the attention effects are

not based on low-level features in the stimulus. Rather, attention effects were observed in areas that are specifically tuned to faces.

References

- Andrews TJ, Davies-Thompson J, Kingstone A, Young AW (2010) Internal and external features of the face are represented holistically in face-selective regions of visual cortex. *J Neurosci* 30:3544–3552. [CrossRef Medline](#)
- Bartels A, Logothetis NK, Moutoussis K (2008) fMRI and its interpretations: an illustration on directional selectivity in area V5/MT. *Trends Neurosci* 31:444–453. [CrossRef Medline](#)
- Beale JM, Keil FC (1995) Categorical effects in the perception of faces. *Cognition* 57:217–239. [CrossRef Medline](#)
- Bland JM, Kerry SM (1998) Statistics notes. Weighted comparison of means. *BMJ* 316:129. [CrossRef Medline](#)
- Brefczynski JA, DeYoe EA (1999) A physiological correlate of the 'spotlight' of visual attention. *Nat Neurosci* 2:370–374. [CrossRef Medline](#)
- Chen AJ, Britton M, Turner GR, Vytlačil J, Thompson TW, D'Esposito M (2012) Goal-directed attention alters the tuning of object-based representations in extrastriate cortex. *Front Hum Neurosci* 6:187. [Medline](#)
- Cohen Kadosh K, Henson RN, Cohen Kadosh R, Johnson MH, Dick F (2010) Task-dependent activation of face-sensitive cortex: an fMRI adaptation study. *J Cogn Neurosci* 22:903–917. [CrossRef Medline](#)
- Connor CE, Gallant JL, Preddie DC, Van Essen DC (1996) Responses in area V4 depend on the spatial relationship between stimulus and attention. *J Neurophysiol* 75:1306–1308. [Medline](#)
- Corbetta M, Miezin FM, Dobmeyer S, Shulman GL, Petersen SE (1991) Selective and divided attention during visual discriminations of shape, color, and speed: functional anatomy by positron emission tomography. *J Neurosci* 11:2383–2402. [Medline](#)
- Cox RW (1996) AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res* 29:162–173. [CrossRef Medline](#)
- Damasio AR, Tranel D, Damasio H (1990) Face agnosia and the neural substrates of memory. *Annu Rev Neurosci* 13:89–109. [CrossRef Medline](#)
- David SV, Hayden BY, Mazer JA, Gallant JL (2008) Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. *Neuron* 59:509–521. [CrossRef Medline](#)
- Desikan RS, Ségonne F, Fischl B, Quinn BT, Dickerson BC, Blacker D, Buckner RL, Dale AM, Maguire RP, Hyman BT, Albert MS, Killiany RJ (2006) An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage* 31:968–980. [CrossRef Medline](#)
- Fairhall SL, Ishai A (2007) Effective connectivity within the distributed cortical network for face perception. *Cereb Cortex* 17:2400–2406. [Medline](#)
- Freiwald WA, Tsao DY, Livingstone MS (2009) A face feature space in the macaque temporal lobe. *Nat Neurosci* 12:1187–1196. [CrossRef Medline](#)
- Gandhi SP, Heeger DJ, Boynton GM (1999) Spatial attention affects brain activity in human primary visual cortex. *Proc Natl Acad Sci U S A* 96:3314–3319. [CrossRef Medline](#)
- Gazzaley A, Cooney JW, McEvoy K, Knight RT, D'Esposito M (2005) Top-down enhancement and suppression of the magnitude and speed of neural activity. *J Cogn Neurosci* 17:507–517. [CrossRef Medline](#)
- Gilaie-Dotan S, Malach R (2007) Sub-exemplar shape tuning in human face-related areas. *Cereb Cortex* 17:325–338. [Medline](#)
- Gross CG (2011) Face neurons. In: *The Oxford handbook of face perception* (Calder AJ, Rhodes G, Johnson MH, Haxby JV, eds), pp 51–76. Oxford: Oxford UP.
- Gross CG, Rocha-Miranda CE, Bender DB (1972) Visual properties of neurons in inferotemporal cortex of the Macaque. *J Neurophysiol* 35:96–111. [Medline](#)
- Harris A, Aguirre GK (2008) The representation of parts and wholes in face-selective cortex. *J Cogn Neurosci* 20:863–878. [CrossRef Medline](#)
- Haxby JV, Hoffman EA, Gobbini MI (2000) The distributed human neural system for face perception. *Trends Cogn Sci* 4:223–233. [CrossRef Medline](#)
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293:2425–2430. [CrossRef Medline](#)
- Kanwisher N, McDermott J, Chun MM (1997) The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J Neurosci* 17:4302–4311. [Medline](#)

- Kay KN, Naselaris T, Prenger RJ, Gallant JL (2008) Identifying natural images from human brain activity. *Nature* 452:352–355. [CrossRef Medline](#)
- Kriegeskorte N, Formisano E, Singer B, Goebel R (2007) Individual faces elicit distinct response patterns in human anterior temporal cortex. *Proc Natl Acad Sci U S A* 104:20600–20605. [CrossRef Medline](#)
- Leopold DA, Bondar IV, Giese MA (2006) Norm-based face encoding by single neurons in the monkey inferotemporal cortex. *Nature* 442:572–575. [CrossRef Medline](#)
- Löffler G, Yourganov G, Wilkinson F, Wilson HR (2005) fMRI evidence for the neural representation of faces. *Nat Neurosci* 8:1386–1390. [CrossRef Medline](#)
- Martínez-Trujillo JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr Biol* 14:744–751. [CrossRef Medline](#)
- McAdams CJ, Maunsell JH (1999) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J Neurosci* 19:431–441. [Medline](#)
- McAdams CJ, Maunsell JH (2000) Attention to both space and feature modulates neuronal responses in macaque area V4. *J Neurophysiol* 83:1751–1755. [Medline](#)
- McCarthy G, Puce A, Gore JC, Allison T (1997) Face-specific processing in the human fusiform gyrus. *J Cogn Neurosci* 9:605–610. [CrossRef](#)
- Minear M, Park DC (2004) A lifespan database of adult facial stimuli. *Behav Res Methods Instrum Comput* 36:630–633. [CrossRef Medline](#)
- Moeller S, Freiwald WA, Tsao DY (2008) Patches with links: a unified system for processing faces in the macaque temporal lobe. *Science* 320:1355–1359. [CrossRef Medline](#)
- Motter BC (1994) Neural correlates of attentive selection for color or luminance in extrastriate area V4. *J Neurosci* 14:2178–2189. [Medline](#)
- Mur M, Ruff DA, Bodurka J, Bandettini PA, Kriegeskorte N (2010) Face-identity change activation outside the face system: “release from adaptation” may not always indicate neuronal selectivity. *Cereb Cortex* 20:2027–2042. [CrossRef Medline](#)
- Murray SO, Wojciulik E (2004) Attention increases neural selectivity in the human lateral occipital complex. *Nat Neurosci* 7:70–74. [CrossRef Medline](#)
- Natu VS, Jiang F, Narvekar A, Keshvari S, Blanz V, O’Toole AJ (2010) Dissociable neural patterns of facial identity across changes in viewpoint. *J Cogn Neurosci* 22:1570–1582. [CrossRef Medline](#)
- Nestor A, Plaut DC, Behrmann M (2011) Unraveling the distributed neural code of facial identity through spatiotemporal pattern analysis. *Proc Natl Acad Sci U S A* 108:9998–10003. [CrossRef Medline](#)
- Pourtois G, Schwartz S, Seghier ML, Lazeyras F, Vuilleumier P (2005) Portraits or people? Distinct representations of face identity in the human visual cortex. *J Cogn Neurosci* 17:1043–1057. [CrossRef Medline](#)
- Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435:1102–1107. [CrossRef Medline](#)
- R Development Core Team (2011) R: A language and environment for statistical computing. Vienna: R Foundation for Statistical Computing.
- Reddy L, Kanwisher NG, VanRullen R (2009) Attention and biased competition in multi-voxel object representations. *Proc Natl Acad Sci U S A* 106:21447–21452. [CrossRef Medline](#)
- Rotshtein P, Henson RN, Treves A, Driver J, Dolan RJ (2005) Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. *Nat Neurosci* 8:107–113. [CrossRef Medline](#)
- Saproo S, Serences JT (2010) Spatial attention improves the quality of population codes in human visual cortex. *J Neurophysiol* 104:885–895. [CrossRef Medline](#)
- Sawamura H, Orban GA, Vogels R (2006) Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the fMRI adaptation paradigm. *Neuron* 49:307–318. [CrossRef Medline](#)
- Serences JT, Saproo S, Scolari M, Ho T, Muftuler LT (2009) Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. *Neuroimage* 44:223–231. [CrossRef Medline](#)
- Somers DC, Dale AM, Seiffert AE, Tootell RB (1999) Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proc Natl Acad Sci U S A* 96:1663–1668. [CrossRef Medline](#)
- Tootell RB, Hadjikhani N, Hall EK, Marrett S, Vanduffel W, Vaughan JT, Dale AM (1998) The retinotopy of visual spatial attention. *Neuron* 21:1409–1422. [CrossRef Medline](#)
- Treue S, Martínez Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399:575–579. [CrossRef Medline](#)
- Treue S, Maunsell JH (1999) Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. *J Neurosci* 19:7591–7602. [Medline](#)
- Valentine T (1991) A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Q J Exp Psychol A* 43:161–204. [Medline](#)
- Winston JS, Henson RN, Fine-Goulden MR, Dolan RJ (2004) fMRI-adaptation reveals dissociable neural representations of identity and expression in face perception. *J Neurophysiol* 92:1830–1839. [CrossRef Medline](#)
- Wojciulik E, Kanwisher N, Driver J (1998) Covert visual attention modulates face-specific activity in the human fusiform gyrus: fMRI study. *J Neurophysiol* 79:1574–1578. [Medline](#)
- Womelsdorf T, Anton-Erxleben K, Pieper F, Treue S (2006) Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. *Nat Neurosci* 9:1156–1160. [CrossRef Medline](#)
- Zhang H, Tian J, Liu J, Li J, Lee K (2009) Intrinsically organized network for face perception during the resting state. *Neurosci Lett* 454:1–5. [CrossRef Medline](#)