Face-Likeness and Image Variability Drive Responses in Human Face-Selective Ventral Regions

Nicolas Davidenko,[1]* David A. Remus,[1] and Kalanit Grill-Spector[1,2]

[1]*Psychology Department, Stanford University, Stanford, California*
[2]*Neuroscience Institute, Stanford University, Stanford, California*

Abstract: The human ventral visual stream contains regions that respond selectively to faces over objects. However, it is unknown whether responses in these regions correlate with how face-like stimuli appear. Here, we use parameterized face silhouettes to manipulate the perceived face-likeness of stimuli and measure responses in face- and object-selective ventral regions with high-resolution fMRI. We first use "concentric hyper-sphere" (CH) sampling to define face silhouettes at different distances from the prototype face. Observers rate the stimuli as progressively more face-like the closer they are to the prototype face. Paradoxically, responses in both face- and object-selective regions decrease as face-likeness ratings increase. Because CH sampling produces blocks of stimuli whose variability is negatively correlated with face-likeness, this effect may be driven by more adaptation during high face-likeness (low-variability) blocks than during low face-likeness (high-variability) blocks. We tested this hypothesis by measuring responses to matched-variability (MV) blocks of stimuli with similar face-likeness ratings as with CH sampling. Critically, under MV sampling, we find a face-specific effect: responses in face-selective regions gradually increase with perceived face-likeness, but responses in object-selective regions are unchanged. Our studies provide novel evidence that face-selective responses correlate with the perceived face-likeness of stimuli, but this effect is revealed only when image variability is controlled across conditions. Finally, our data show that variability is a powerful factor that drives responses across the ventral stream. This indicates that controlling variability across conditions should be a critical tool in future neuroimaging studies of face and object representation. *Hum Brain Mapp 33:2334–2349, 2012.*

Key words: face-likeness; variability; face silhouettes; FFA; fMRI

INTRODUCTION

Functional magnetic resonance imaging (fMRI) studies reliably identify regions in the human ventral stream that respond selectively to faces [Kanwisher et al., 1997; Puce et al., 1995; Tong et al., 2000; Yovel and Kanwisher, 2004], including regions in the middle (*mFus-faces*) and posterior [*pFus-faces* (Weiner and Grill-Spector, 2010)] fusiform gyrus (often referred to collectively as FFA), and a region in lateral occipital cortex, overlapping the inferior occipital gyrus, referred to as OFA [Gauthier et al., 2000], or *IOG-faces* [Weiner and Grill-Spector, 2010]. These face-selective regions are thought to be critically involved in face perception [Andrews et al., 2002; Hasson et al., 2001; Moutoussis and Zeki, 2002; Rotshtein et al., 2005; Tong et al., 1998] and recognition [Barton, 2002; Golarai et al., 2007; Grill-

Additional Supporting Information may be found in the online version of this article.

Contract grant sponsor: NIH NRSA; Contract grant number: EY18279-01A1; Contract grant sponsor: Whitehall; Contract grant number: 2005-05-111-RES; Contract grant sponsor: NSF; Contract grant numbers: BCS 0920865, 0547775, 0624345; Contract grant sponsor: NEI; Contract grant number: RO1EY019279-01-A1; Contract grant sponsor: American Philosophical Society.

*Correspondence to: Nicolas Davidenko, Psychology Department, Stanford University, Stanford, California.
E-mail: ndaviden@gmail.com

Received for publication 21 December 2010; Revised 14 April 2011; Accepted 2 May 2011
DOI: 10.1002/hbm.21367
Published online 5 August 2011 in Wiley Online Library (wileyonlinelibrary.com).

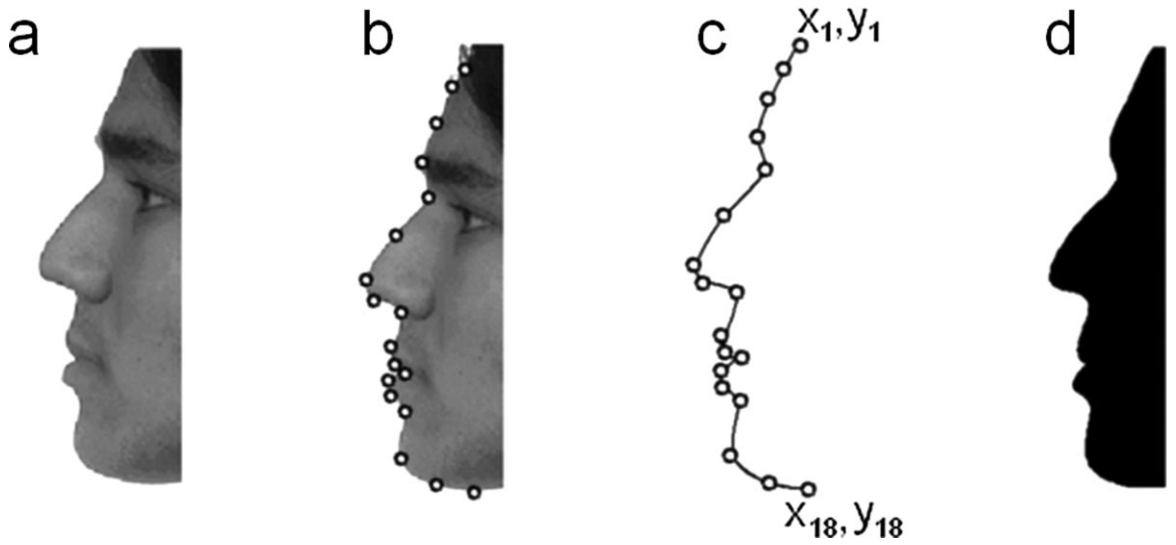

Figure 1.
Steps in parameterizing a face silhouette [adapted from Davidenko, 2007]. (**a**) A profile face image from the FERET database; (**b**) 18 keypoints are identified along the profile contour; (**c**) keypoints are normalized into a common *x*–*y* plane across all face images; (**d**) bi-cubic splines smoothly interpolate between keypoints to create the parameterized face silhouette.

Spector et al., 2004], and much work in the past decade has focused on understanding their response properties.

To date it has not been established how face-selective regions respond to stimuli that vary in their degree of face-likeness. On the one hand, responses in face-selective regions are higher when stimuli are perceived as faces vs. when they are not. For example, when presented with ambiguous stimuli (e.g., the Rubin face-vase illusion) responses in fusiform face-selective regions are higher when subjects report seeing two faces rather than a vase [Andrews et al., 2002; Hasson et al., 2001]. Further, several studies report that face-selective regions respond similarly to different kinds of faces, such as familiar vs. unfamiliar faces [Ewbank and Andrews, 2008; Kanwisher et al., 1997], adult vs. child faces [Golarai et al., 2010], male vs. female faces [Kranz and Ishai, 2006], and front- vs. profile-view faces [Kanwisher et al., 1997] (though there is evidence for some modulation of response by race [Golby et al., 2001] and distinctiveness [Loffler et al., 2005]). On the other hand, although face-selective regions respond more to faces than non-faces, their responses are not uniformly low across non-face stimuli. For example, face-selective regions respond more strongly to images of animals and body parts than to inanimate objects, and more strongly to objects than to houses and scenes [Downing et al., 2006; Grill-Spector, 2003; Grill-Spector et al., 2004; Weiner and Grill-Spector, 2010]. Because these visual categories differ on many dimensions, it is unclear what aspects of images modulate the level of response of face-selective regions. A recent study reported that FFA responses to images at various morph levels between a face and a house decreased monotonically as the morph level had a lower contribution of the face image, suggesting that responses may be modulated by stimulus shape [Tootell, 2008]. However, several intermediate morphs between the face and the house resembled neither a face nor a house, and there were no behavioral measurements of how face-like the stimuli appeared to subjects. Thus, it remains unknown what kinds of stimulus transformations modulate response levels in face-selective regions, and how they might relate to perceived face-likeness.

To address these questions we used a novel method of parameterized face silhouettes [Davidenko, 2007]. This face space is built from a database of 480 profile-view, gray-scale faces which are parameterized into silhouettes using 18 keypoints (see Fig. 1 and Materials and Methods). The resulting face stimuli are defined entirely by their shape. Using principal component analysis (PCA), we capture the shape dimensions, or principal components (PCs), along which human face profiles vary relative to a prototype face (the average of these 480 silhouettes, which subjects perceive as typical [Davidenko, 2007; Davidenko and Ramscar, 2006; Davidenko et al., 2007]). Importantly, the face silhouette methodology provides a method for manipulating the perceived face-likeness of stimuli simply by transforming the shape of silhouettes along the PCs that faces typically vary along, while controlling low-level properties of the stimuli, such as size, contrast, and spatial frequency.

Here, we generated silhouettes at various distances from the prototype face. We first measured the perceived face-likeness of these stimuli and show that faces further away

from the prototype are perceived as less face-like, in a monotonically decreasing way. We then measured responses in face- and object-selective ventral regions with high-resolution fMRI. We reasoned that responses in face-selective, but not object-selective, regions would be modulated by the degree of face-likeness of the silhouettes. We predicted that responses in fusiform face-selective regions should decrease as face-like ratings decrease, since responses in these regions are known to be correlated with subjects' face perception.

MATERIALS AND METHODS

Subjects

Thirteen right-handed subjects (seven female) participated in one or more of four fMRI experiments. In Experiments 1 and 2 (Face silhouette study), we verified that face silhouettes elicit a face-selective response profile in human ventral visual cortex. Eight subjects (six female, ages 23–39) participated in Experiment 1 and 9 subjects participated in Experiment 2. In Experiments 3 and 4, we manipulated face-likeness and measured responses in face- and object-selective regions. Ten right-handed subjects (five female, ages 21–39 years) participated in Experiment 3 (CH sampling) and nine right-handed subjects (four female, ages 21–39) participated in Experiment 4 (MV sampling); seven subjects participated in both of these. Stanford University's Human Subjects IRB approved the experimental protocol and subjects gave written informed consent to participate.

fMRI Scans

Scans were conducted at the Richard M. Lucas Center for Imaging at Stanford University using an eight-channel surface coil in a 3-Tesla GE magnet. Experiments 1–4 were conducted on separate days between 3 and 11 months apart for different subjects (mean: 6 months). Each scan lasted ~90 min and included an anatomical run, two functional localizer runs to identify face- and object-selective regions, and eight experimental runs. Functional scans were conducted at high-resolution (HR-fMRI, 1.5 mm isotropic voxels) and covered a field of view of 19.2 cm × 19.2 cm × 3.9 cm that included occipital and temporal regions. HR-fMRI allows for better segregation between object- and face-selective activations and more accurate ROI definitions. We acquired 26 1.5 mm-thick oblique slices using a two-shot T2*-sensitive spiral acquisition sequence [Glover, 1999] (TE = 30 ms, TR = 2,000 ms, flip angle = 77°, and bandwidth = 125 kHz). In-plane T1-weighted anatomical images were acquired in the same session as the functional scans using the same prescription as the functional scans and a standard two-dimensional RF-spoiled GRASS (SPGR) pulse sequence (TE = 1.9 ms, flip angle = 15°, bandwidth = 15.63 kHz). The anatomical in-planes were used to co-register each subject's functional and anatomical data to aid with ROI visualization and definition. A high-resolution anatomical volume of the whole brain was acquired for 11 subjects with a head-coil using a T1-weighted SPGR pulse sequence (TR = 1,000 ms, FA = 45°, 2 NEX, FOV = 200 mm, resolution of 0.78 × 0.78 × 1.2 mm^3). For the remaining two subjects, ROI definitions were conducted based on the inplane images acquired in each session.

fMRI Data Analysis

Imaging data was analyzed using MATLAB and our in-house software, *mrVista* (white.stanford.edu/software). The analysis included alignment of functional data to whole-brain anatomies, motion correction of the functional data, transformation of each voxel's time course to percent signal change, and application of a general linear model (GLM [Worsley et al., 1996] using the hemodynamic response function as implemented in SPM2 and SPM5, www.fil.ion.ucl.ac.uk/spm] to estimate β-weights for each voxel in each experimental condition. There was no spatial smoothing. For ROI-based analyses, we measured the average in each subject's individually defined ROIs from the localizer scan from each condition of Experiments 1–4. We then averaged data across subjects. Data were analyzed separately for each hemisphere.

Localizer Runs

Two 5-min localizer scans were conducted at the beginning of each scanning session. Subjects observed 12-s blocks of grayscale front-view faces, grayscale common objects (cars, instruments, and plants), gray-scale scrambled versions of these images, and fixation blocks, with 12 instances of each type of block using different stimuli. These localizer runs were analyzed to define face- and object-selective ROIs separately in each subject.

Face-selective ROIs were defined in each subject as voxels in the fusiform gyrus (Fus) and lateral occipital cortex overlapping the inferior occipital gyrus (IOG) that responded significantly more strongly to faces than to objects ($P < 0.001$, voxel level, uncorrected; see Fig. 2–red). In most subjects, right hemisphere Fus-faces consisted of a posterior (pFus-faces) and anterior (mFus-faces) region that we analyzed separately. The mFus-faces ROI is more medial and tends to overlap the mid-fusiform sulcus, whereas the pFus-ROI is about 1–2 cm more posterior and sometimes overlaps the occipito-temporal sulculs (OTS), see [Weiner and Grill-Spector, 2010]. We focus our analyses on right- and left-hemisphere ROIs that were present in at least five subjects in each of the experiments. These include right-hemisphere mFus-faces (minimum $n = 9$), pFus-faces (9), and IOG-faces (6) and left-hemisphere pFus-faces (9). Left mFus-faces (minimum $n = 3$) and IOG faces (4) were found less consistently in our studies.

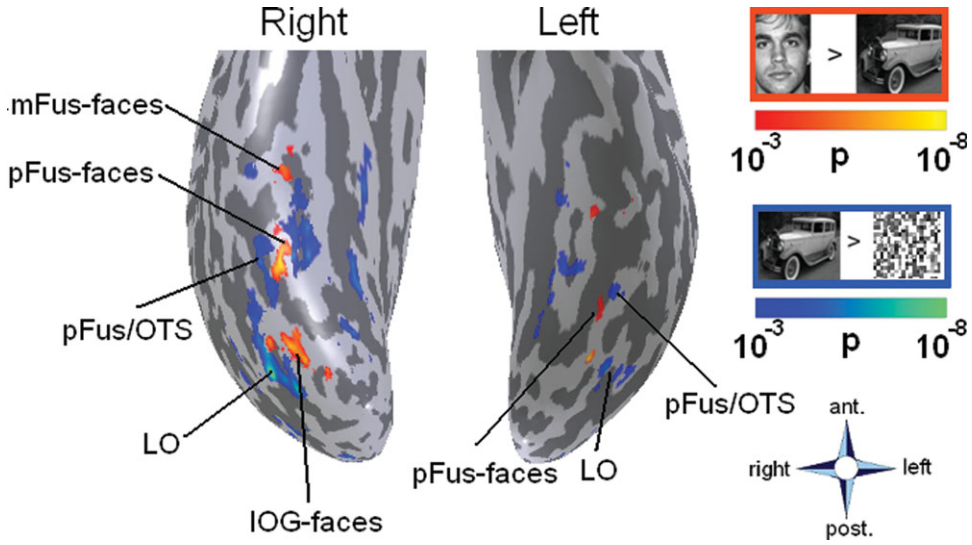

Figure 2.

Face- and object-selective activations on a ventral surface of a representative subject's inflated cortical surface. Face-selective regions (responding more to front-view faces than objects) are shown in red and object-selective regions (responding more to objects than scrambled images) are shown in blue. Color bar indicates the statistical significance (*P*-value) of the contrast of interest. Abbreviations: mFus: mid-fusiform; pFus: posterior fusiform: OTS: occipito-temporal sulcus; IOG: inferior occipital gyrus. LO: lateral occipital object-selective region.

Object-selective regions were defined as voxels in posterior fusiform gyrus overlapping with the OTS (pFus/OTS) and in the lateral aspect of the occipital lobe posterior and ventral to MT (LO) that responded significantly more strongly to objects than scrambled images ($P < 0.001$, voxel level, uncorrected; Fig. 2–blue). We identified object-selective regions in the right hemisphere pFus/OTS (minimum $n = 6$) and LO (9) and left hemisphere pFus/OTS (6) and LO (7).

Results of Experiments 1–4 are shown as the average percentage signal for each condition across participants, shown separately in each face- and object-selective ROI. For robustness and to complement our ROI analyses, we measured the extent of activation within each region to the different experimental conditions. We defined a voxel count measure in Experiments 3 and 4 as the number of voxels within each functional ROI responding significantly ($P < 0.01$) more strongly to each active condition of Experiments 3 and 4 versus fixation. This analysis was done separately for each subject and ROI. Average voxel counts across subjects are shown in Supporting Information Figures 5 and 6.

Stimuli

Our stimuli were constructed using the parameterized face silhouettes methodology [Davidenko, 2007]. Silhouette face space is generated by placing 18 keypoints along the contour of 480 left-facing profile photographs of male and female faces from the FERET face database [Phillips, 1998, 2000] (Fig. 1a). We fixed the top and bottom points to control for stimulus size and orientation. The XY coordinates of the remaining 16 keypoints form the basis of a principal components analysis (PCA) that produces 32 orthogonal dimensions, or principal components (PCs) that fully describe the space of face silhouettes [Davidenko, 2007]. The prototype face silhouette—the average of all 480 face profiles—is obtained by setting 0-coefficients along all PCs and novel face silhouettes can be generated by sampling the PC coefficients from a multi-normal distribution [see Davidenko, 2007].

In previous work [Davidenko, 2007], we have shown face silhouettes are perceived much like regular face stimuli, eliciting accurate judgments of gender and age, reliable ratings of attractiveness and distinctiveness, successful identification with their front-view, gray-scale counterparts, and better recognition when upright than upside-down. In Experiments 1 and 2, we further show that parameterized face silhouettes elicit a face-selective response profile in the ventral stream.

Face Silhouette Study (Experiments 1 and 2)

To determine whether face silhouettes elicit a face-selective response profile, we measured responses with high-resolution fMRI as subjects observed blocks of face silhouettes, two-tone shapes, and scrambled two-tone images (Experiment 1), and blocks of upright and upside-

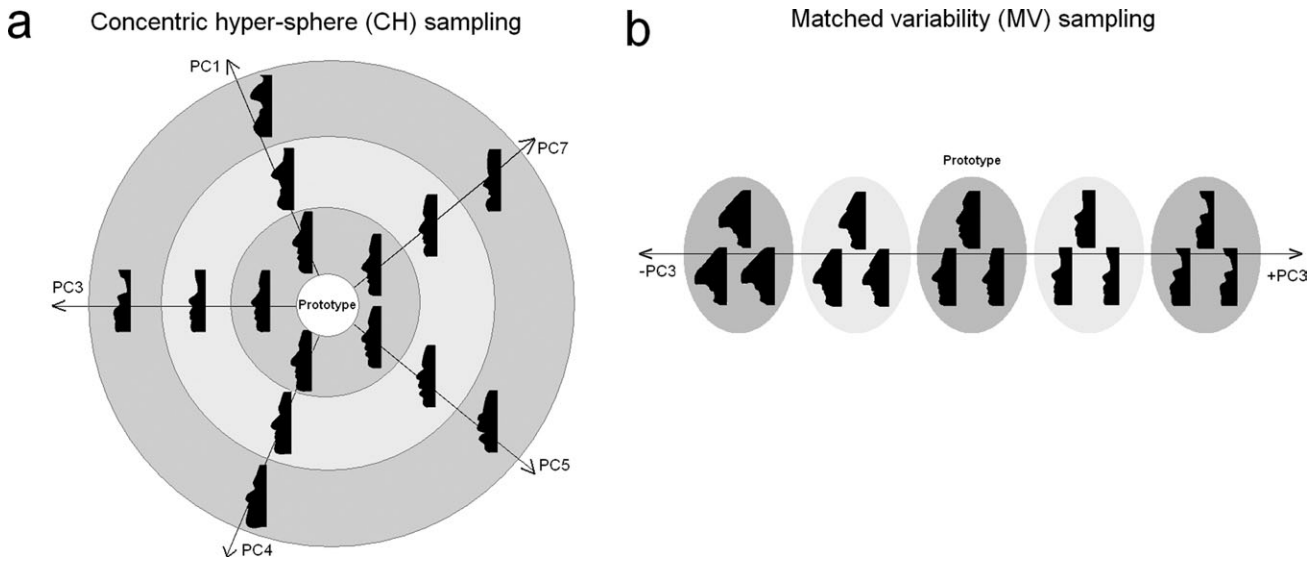

Figure 3.

(**a**) Schematics of CH and MV sampling. Each line represents a different PC and each shaded disk represents a block of stimuli at a given distance from the prototype to be shown in the fMRI experiment. (**b**) Schematic diagram of MV sampling for Experiment 4. Silhouettes are defined along PC3 and each equal-size shaded disk represents a block of matched-variability stimuli at a given distance from the prototype to be shown in the fMRI experiment.

down face silhouettes (Experiment 2). In Experiment 1, we generated multiple face silhouettes by sampling the coefficients of the first 6 PCs from a multi-normal distribution around the prototype face. Each PC coefficient was weighted by its *z*-score, thus producing normally distributed values along each dimension, resulting in realistic looking face silhouettes. We used six dimensions in order to generate a wide variety of silhouette stimuli that would not repeat across blocks and runs. Using a similar keypoint construction method, we generated two-tone shapes that did not resemble faces but were matched with face silhouettes in their size, contrast, and number of keypoints. We generated scrambled two-tone images by splitting the face and shape silhouette images into 8 by 8 subimages and randomizing the positions of the sub-images. In Experiment 2, we generated upright face silhouettes by varying the coefficients of PCs 2 and 6 by up to ± 1.5 and ± 2 standard deviations (PC units), respectively, away from the prototype face silhouette and upside-down silhouettes by vertically inverting the upright silhouettes.

For each subject in Experiments 1 and 2, we acquired 8 runs where subjects observed six 12-s image blocks alternating with 12-s fixation blocks. Each run included one instance of each type of image block. Stimuli were displayed in MATLAB (Mathworks.com) using Psychtoolbox2 [Brainard, 1997] and projected onto a screen inside the scanner. Images subtended $\sim 5.7^\circ$ of subjects' visual field and were presented at 1 Hz, for 750 ms with 250 ms fixation interstimulus intervals. During all image blocks, subjects fixated on a central cross and performed a one-back task, pressing a button whenever two consecutive images repeated ($\sim$8% of trials). Subjects' performance was monitored to ensure they were paying attention to the stimuli.

Measuring Responses in the Ventral Stream as a Function of Face-Likeness

Critically, we contrast two methods of defining blocks of silhouettes at different levels of face-likeness. Concentric hyper-sphere (CH) sampling (Fig. 3a and Supporting Information Fig. 1) is a method that has been used previously to manipulate face distinctiveness [Loffler et al., 2005; Valentine, 1991]. CH sampling captures a wide range of face stimuli but results in a correlation between distances from the prototype face and image variability, whereby blocks of stimuli far from the prototype are more variable than blocks of stimuli near the prototype. To account for this potential confound, we introduce a novel method of "matched variability" (MV) sampling, where silhouettes are defined primarily along one PC of face space at the same distances from the prototype, but with each block of silhouettes spanning an equal-size region of face space and thus matching variability across blocks (Fig. 3b).

Measures of Variability

We defined a metric of image variability to describe the heterogeneity of images in each block. Image variability was defined as the mean distance in face space between

pairs of 32-dimensional vectors corresponding to pairs of face silhouettes within each block. The units of image variability are in standard deviations of the PC coefficients (or PC units). This measure was nearly identical to a measure based on the sum of XY distances between pairs of corresponding keypoints that define each silhouette stimulus ($r = 0.9993$, $P < 10^{-10}$). We considered two additional metrics of variability in our analyses: pixel-based variability and perceptual variability. Pixel-based variability was defined as the pixel-by-pixel difference between the two-tone images in each block, expressed as a proportion of different pixels in the image area. Perceptual variability was measured in a post-scan study outside the scanner. Nine of our subjects provided ratings of dissimilarity among pairs of silhouettes sampled from each of the blocks, on a seven-point scale, with "1" representing "identical" and "7" representing "maximally dissimilar."

Concentric Hyper-Sphere (CH) Sampling (Experiment 3)

We generated silhouettes by sampling from 12 orthogonal face space dimensions (PCs 1, 3–5, and 7–14) at 9 different distances (0, 1.5, 3, 4.5, 6, 7.5, 9, 10.5, and 12 PC units) from the prototype (see Fig. 3a), and blocking the silhouettes by their distance from the prototype. As with image variability, distances from the prototype are measured along each PC as a *z*-score (number of standard deviations) of that PCs coefficients in the silhouette face space parameterization [Davidenko, 2007] and we refer to these as PC units. The coefficients of PCs 2 and 6 also varied by up to ±1.5 PC and ±2 PC units, respectively, to generate different exemplar silhouettes across blocks and runs. We selected PCs 2 and 6 for this because they were found to be least correlated with social attributes of faces such as gender and race [Davidenko and Ramscar, 2006]. Each block thus contained stimuli sampled from 12 different dimensions of face space, covering a broad range of face silhouettes. Importantly, image variability was not matched across CH blocks. Instead, mean image variability had values of 3.2, 4.0, 5.5, 7.3, 9.2, 11.2, 13.2, 15.3, and 17.3 PC units: as blocks of stimuli deviated more from the prototype face silhouette, within-block image variability increased (see Fig. 3a and Supporting Information Fig. 2).

Matched Variability (MV) Sampling (Experiment 4)

Here, we defined nine different blocks of face silhouettes by sampling along both directions of PC 3, a dimension associated with face gender [Davidenko, 2007], at 0, ±3, ±6, ±9, and ±12 PC units from the prototype (see Fig. 3b). We selected PC 3 because it produced a representative range of face-likeness ratings in a pilot study (see Supporting Information Fig. 3), but we would predict similar results if other dimensions were used. Different exemplars in each block were generated by varying the coefficient of PCs 2 and 6 by up to ±1.5 and ±2 PC units, respectively, as with CH sampling. Because the size of each sampling region was kept constant, image variability was matched at 3.2 PC units across all blocks, equivalent to the lowest variability blocks in CH sampling (Experiment 3).

For each subject in Experiments 3 and 4, we acquired eight runs where subjects observed 9 12-s image blocks alternating with 12-s fixation blocks. Each run included one block at each distance from the prototype face. In each run, the blocks were presented in an order that minimized any cross-block adaptation. Specifically, in Experiment 3 (CH sampling), the general order of blocks (labeled by distance from the prototype face, in PC units) was 0, 7.5, 1.5, 9, 3, 10.5, 4.5, 12, and 6. Across runs, this order was alternately reversed, and the starting block was randomized. Similarly, in Experiment 4 (MV sampling), the general order of blocks was −12, 3, −9, 6, −6, 9, −3, 12, 0, keeping the distance in face-space between consecutive blocks nearly constant to avoid any systematic biases in block-to-block similarity that might differentially affect adaptation in some blocks. The stimulus display method, timing, and one-back task were identical to those in Experiments 1 and 2.

Perceived Face-Likeness

Following the fMRI scans, nine of the subjects were given a questionnaire that included 54 face silhouettes (three from each of the MV and CH blocks) as well as three additional two-tone shape stimuli that matched the face silhouettes in size, contrast, and number of keypoints but did not resemble faces (Fig. 4c). Subjects provided a 1–5 face-likeness rating, where 1 indicated "Not at all face-like" and 5 indicated "Completely face-like." As shown in Figure 4a,b, stimuli near the prototype were rated as highly face-like, and ratings decreased gradually with distance from the prototype. Although this was true for all PCs, equal distances along different PCs (or along different directions of the same PC) can yield slightly different subjective ratings of face-likeness. For example, face-likeness ratings decrease faster along the positive direction of PC3 than along the negative direction (Fig. 4b). This suggests that face-likeness ratings may be better than distance from the prototype in characterizing brain responses in relation to face perception. Thus, in subsequent fMRI analyses, responses are plotted as a function of the mean face-likeness rating associated with each silhouette block. Importantly, even though face silhouettes in CH and MV sampling were defined along different PCs, face-like ratings covered a similar range across studies (Fig. 4a,b), from a mean ± SEM rating of 4.9 ± 0.1 for stimuli closest to the prototype to 2.5 ± 0.2 for stimuli farthest from the prototype. We note that the lowest ratings on face silhouettes did not reach the floor of the scale. Only the nonface

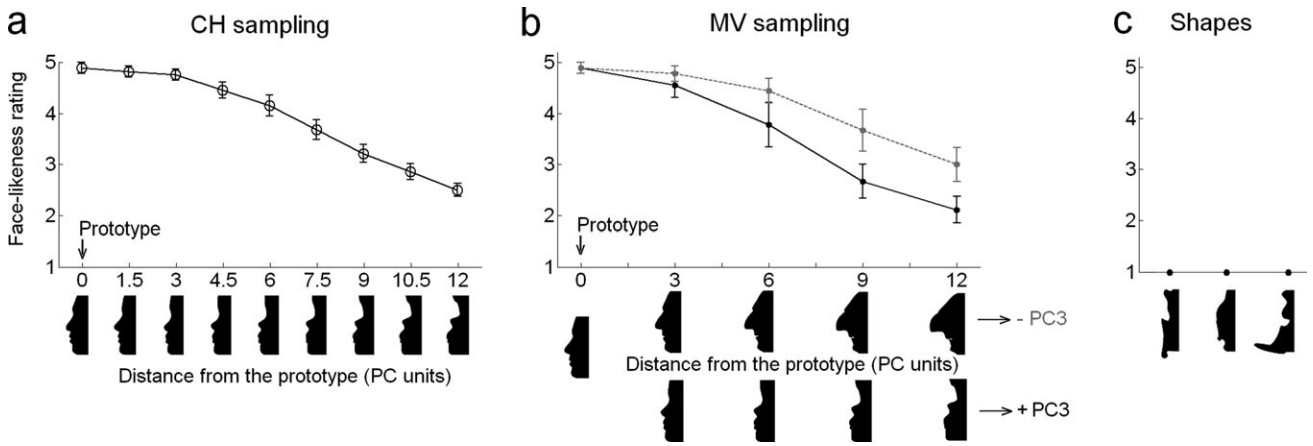

Figure 4.

Face-like ratings decrease as stimuli deviate from the prototype face for both CH and MV sampling. (**a**) Face-like ratings on stimuli from the CH blocks (Experiment 3) provided outside the scanner by nine subjects who participated in the fMRI experiments, plotted as a function of distance from the prototype face, measured in PC units; (**b**) Face-like ratings on stimuli from the MV blocks (in Experiment 4) provided by the same nine subjects. *X*-axis indicates distance from the prototype face along PC3 in the negative (gray) and positive (black) directions. (**c**) Ratings on similarly constructed nonface shapes were uniformly 1 across subjects. Face-likeness scale: 1 = Not at all face-like to 5 = Completely face-like. Error bars denote between-subjects SEM.

shapes received a rating of 1 ("not at all face-like"), and this was observed for every subject (Fig. 4c).

RESULTS

Face Silhouettes Elicit a Face-Selective Response Profile in Face-Selective Regions

To test the utility of face silhouettes as a tool for fMRI investigations of face perception, we first examined whether face silhouettes elicit a face-selective response profile. In Experiment 1 (8 subjects), we compared responses in face- and object-selective regions to face silhouettes, two-tone shapes that were controlled for low-level properties, and scrambled two-tone images (see Materials and Methods). In Experiment 2 (9 subjects), we compared responses to blocks of upright face silhouettes to the same silhouettes presented upside-down.

We extracted the mean response in each image condition from each subject's ROIs defined from the independent localizer (see Materials and Methods) and averaged these responses across subjects. In Experiment 1, all face-selective ROIs responded significantly more strongly during blocks of face silhouettes than during blocks of two-tone shapes (ts > 2.7, Ps < 0.03; paired t-test, two-tailed) and scrambled two-tone images (ts > 4.3, $P < 0.02$, Fig. 5a). The voxels responding significantly more to face versus shape silhouettes ($P < 0.01$) overlapped with the independently localized face-selective ROIs (see Supporting Information Fig. 4).

In Experiment 2, all face-selective ROIs responded significantly more strongly to upright than to upside-down face silhouettes (ts > 3.3, Ps < 0.01, Fig. 5a), demonstrating an fMRI face-inversion effect with face silhouettes. In contrast, responses in object-selective regions (pFus/OTS and LO) were higher for intact vs. scrambled two-tone images (ts > 2.1, Ps < 0.05) but similar across face silhouettes and two-tone shapes in Experiment 1 ($|t$s$| < 1.1$, Ps > 0.3) and between upright and upside-down face silhouettes in Experiment 2 ($|t$s$| < 0.6$, Ps > 0.5; Fig. 5b). These data complement our work showing face silhouettes' effectiveness as stimuli for behavioral tasks [Davidenko, 2007] and previous fMRI studies that have shown higher responses in face-selective regions when subjects report a face percept in Rubin's face/vase stimuli [Andrews et al., 2002; Hasson et al., 2001]. Taken together, these results indicate that face silhouettes elicit a face-selective response profile and are therefore useful stimuli for fMRI studies of face representation.

Concentric Hypersphere Sampling: Responses in Both Face- and Object-Selective are Negatively Correlated With Face-Likeness

We measured responses in face-selective regions, defined from the independent localizer, as subjects observed silhouettes at nine different levels of face-likeness within a five-point scale, ranging from 2.5 (low face-likeness) to 4.9 (high face-likeness; see Fig. 4a and Materials and Methods). To our surprise, we found that average

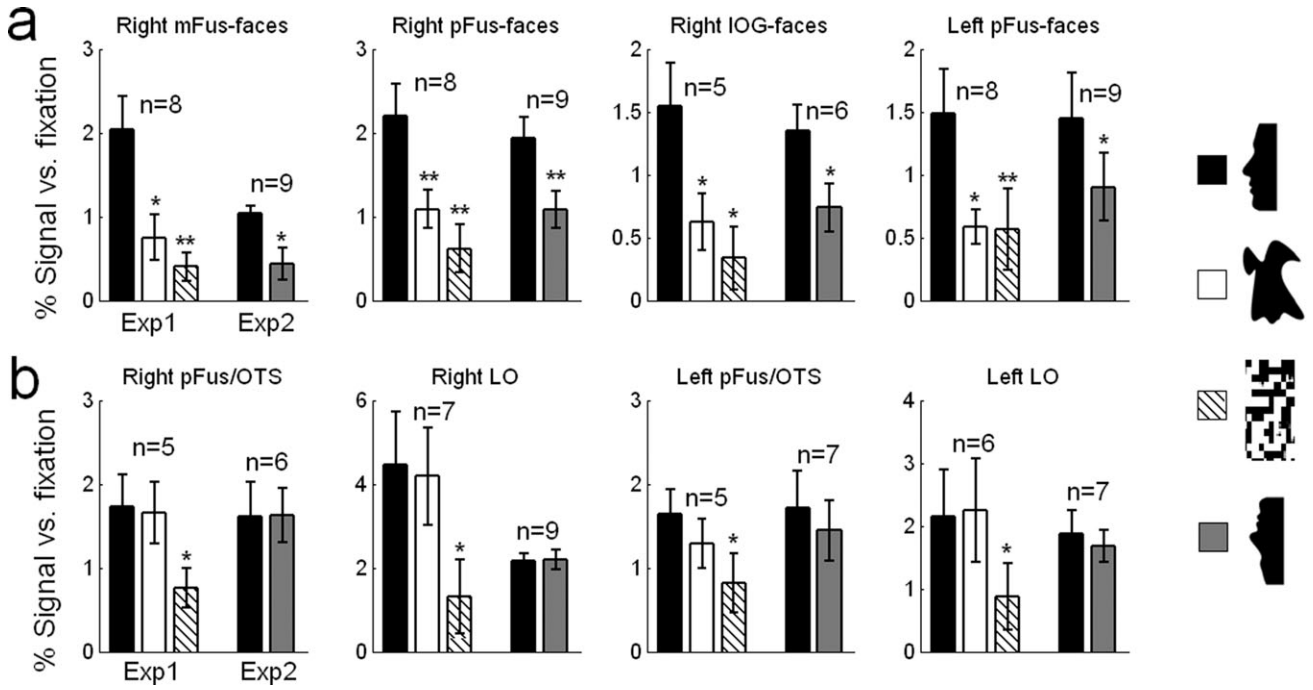

Figure 5.

Responses in face-selective (but not object-selective) regions are higher for face silhouettes than both shape silhouettes and upside-down face silhouettes. Responses in face-selective (**a**) and object-selective ROIs (**b**) to upright face silhouettes (black), two-tone shapes (white), scrambled two-tones (striped), and upside-down silhouettes (gray) in the face silhouette study (Experiments 1 and 2). Ns indicate number of subjects in which each ROI was localized in the independent localizer scan. Asterisks indicate responses significantly different than upright silhouettes (*$P < 0.05$; **$P < 0.001$; paired *t*-test, two-tailed).

responses across subjects in fusiform and IOG face-selective ROIs were significantly negatively correlated with perceived face-likeness ($r_{\text{Right mFus-faces}} = -0.91$, $P < 0.001$; $r_{\text{Right pFus-faces}} = -0.88$, $P < 0.01$; $r_{\text{Right IOG-faces}} = -0.9$, $P < 0.001$; $r_{\text{Left pFus-faces}} = -0.85$, $P < 0.01$; see Fig. 6a). That is, responses in face-selective regions decreased as face-likeness increased. In fact, the strongest responses in right mFus-faces and right pFus-faces were observed during blocks of stimuli rated by the same observers as being the least face-like. Moreover, this puzzling pattern of responses was also present in object-selective regions. Average responses in object-selective regions were also negatively correlated with perceived face-likeness ($r_{\text{Right pFus/OTS}} = -0.94$, $P < 10^{-4}$; $r_{\text{Right LO}} = -0.97$, $P < 10^{-4}$; $r_{\text{Left pFus/OTS}} = -0.85$, $P < 0.01$; $r_{\text{Left LO}} = -0.91$, $P < 0.001$; see Fig. 6b) even though these regions did not show face-selective responses in Experiments 1 and 2. When we calculated the number of voxels within these ROIs activated by each stimulus condition versus fixation ($P < 0.01$), we found similar results: the number of active voxels decreased as perceived face-likeness increased (Ps < 0.01 in every ROI; see Materials and Methods and Supporting Information Fig. 5).

Why are responses in these ventral regions negatively correlated with perceived face-likeness?

We suggest that these counterintuitive results are not driven by the perceived face-likeness of stimuli itself, but by the image variability that differs across stimulus blocks. Research on fMRI-adaptation [Andrews and Ewbank, 2004; Avidan et al., 2002; Grill-Spector and Malach, 2001; Grill-Spector et al., 1999, 2006; Li et al., 1993; Yovel and Kanwisher, 2004] has shown that responses in high-level visual areas, including these face- and object-selective regions, are reduced (adapted) when presented repeatedly with identical stimuli. Further, fMRI-A effects generalize across similar stimuli. In other words, researchers report more fMRI-A to similar stimuli than to dissimilar stimuli, although precisely how the degree of stimulus similarity determines the level of fMRI-A is debated and differs across brain regions [Gilaie-Dotan and Malach, 2007; Gilaie-Dotan et al., 2010; Weiner et al., 2010]. In CH sampling, highly face-like silhouettes are sampled from relatively small regions of face space near the prototype (Fig. 3a), resulting in low variability blocks that may generate considerable fMRI-A. In contrast, low face-likeness blocks further from the prototype are defined in large regions of face space and are more variable, which may lead to less fMRI-A and consequently a stronger signal. Indeed, there is a high negative correlation between mean image variability and mean rated

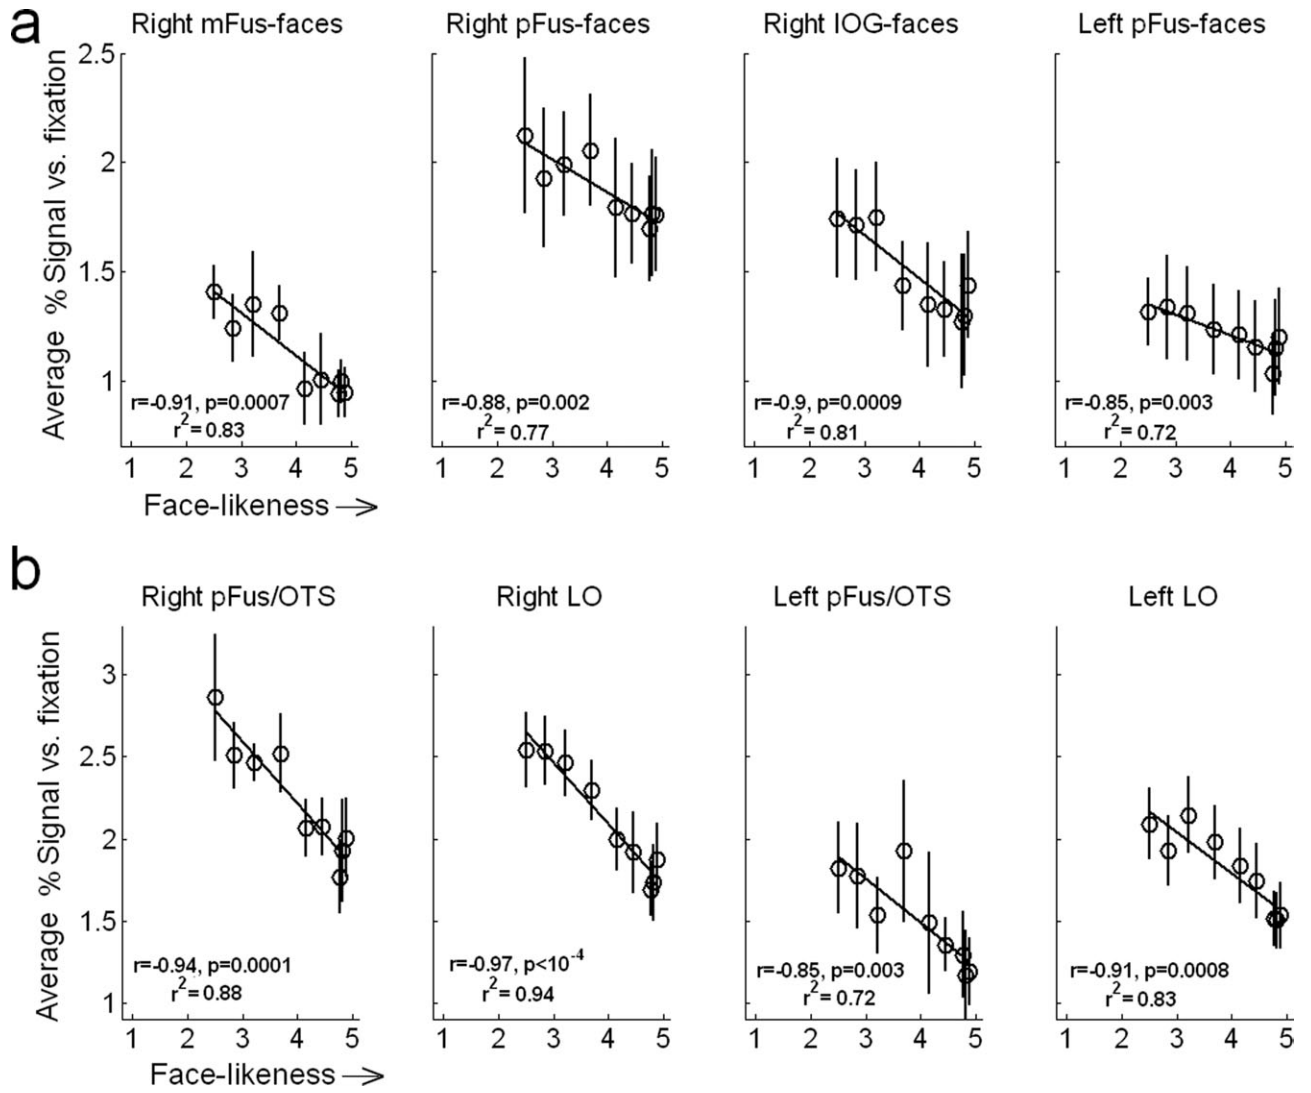

Figure 6.

CH sampling; responses in both face- and object-selective ROIs decrease as stimuli become more face-like. Each point reflects the average response across subjects to a block of stimuli. *r*- and *P*-values denote correlations and their significance, respectively, between average face-likeness ratings and average percentage signal across subjects; error bars denote between-subjects SEM. R^2 indicates proportion variance explained. (**a**) Average percentage signal change across subjects in face-selective ROIs in CH sampling across subjects as a function of perceived face-likeness of stimuli; Number of subjects in which each ROI was localized: Right mFus-faces ($n = 9$), Right pFus-faces ($n = 10$), Right IOG-faces ($n = 8$); Left pFus-faces ($n = 10$); (**b**) Average percentage signal change across subjects in object-selective ROIs in CH sampling. Right pFus/OTS ($n = 8$), Right LO ($n = 10$), Left pFus/OTS ($n = 6$), Left LO ($n = 8$).

face-likeness ratings ($r = -0.99$, $P < 10^{-6}$). Therefore, it is possible that image variability contributed to the differential signal amplitudes across blocks. To test this hypothesis and isolate the effect of face-likeness from effects of variability on responses, we conducted another experiment where we manipulated face-likeness as in CH sampling but kept image variability constant across blocks.

Matched Variability Sampling: Responses in Face- (But Not Object-) Selective Regions are Positively Correlated With Face-Likeness

We measured activity in the same face- and object-selective ROIs in nine subjects (seven overlapping with Experiment 3) while they observed blocks of face silhouettes defined at the same distances from the prototype and with

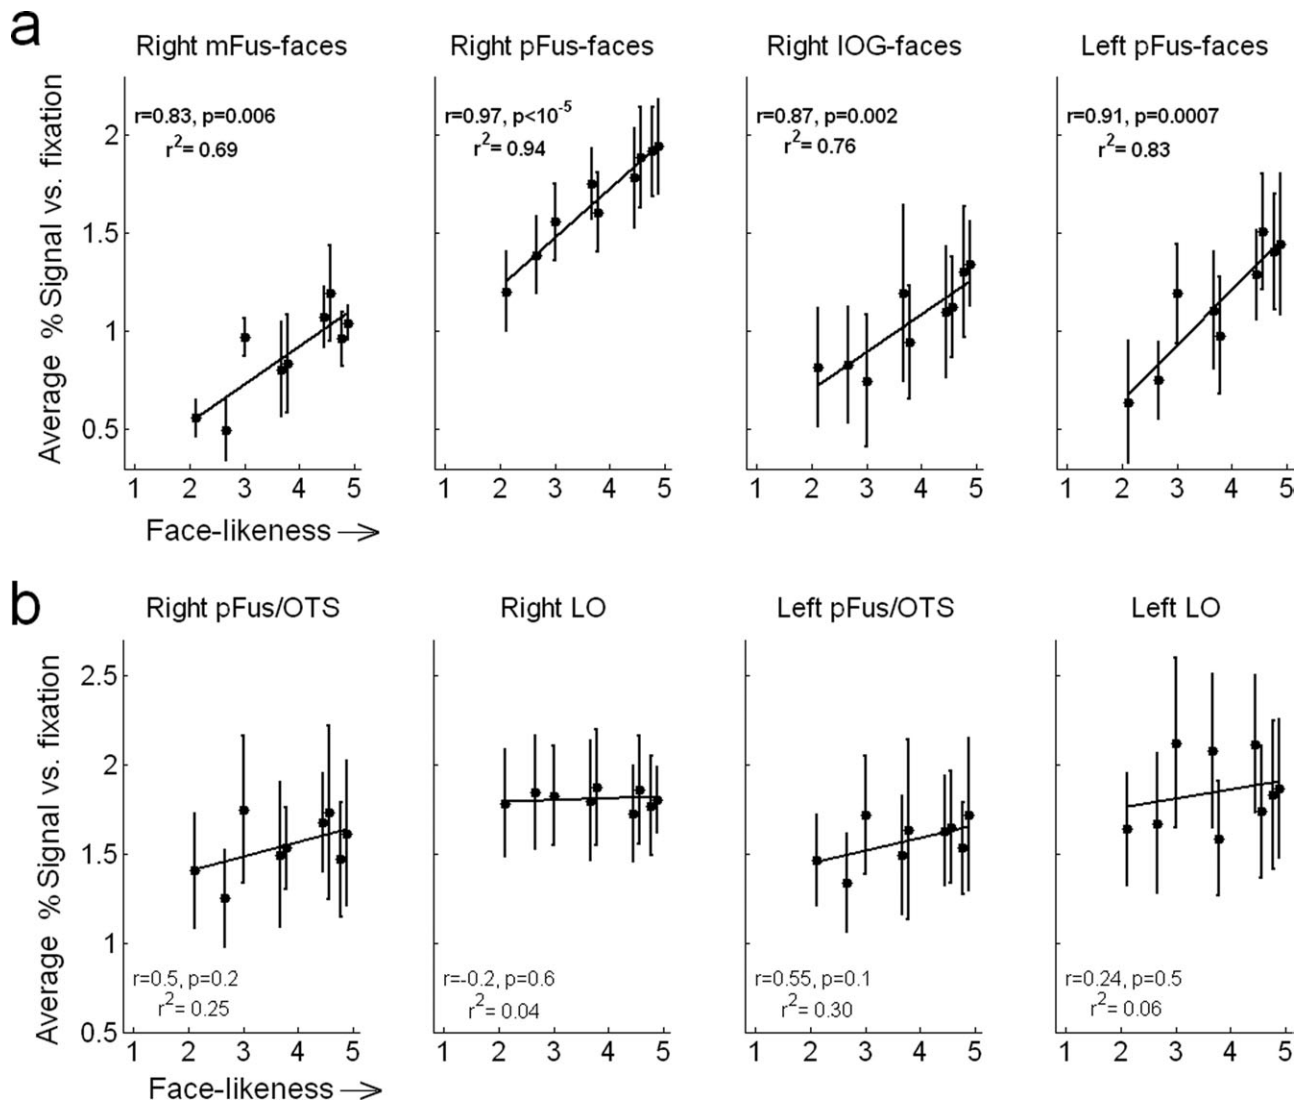

Figure 7.

MV sampling; responses in face-selective (but not object-selective) ROIs increase as stimuli become more face-like. Each point reflects the average response across subjects to a block of stimuli. *r*- and *P*-values denote correlations and their significance, respectively, between face-likeness and average percentage signal across subjects; error bars denote between-subjects SEM. (**a**) Average percentage signal change in face-selective ROIs in MV sampling across subjects as a function of rated face-likeness of stimuli; Number of subjects in which each ROI was localized: Right mFus-faces ($n = 9$), Right pFus-faces ($n = 9$), Right IOG-faces ($n = 6$); Left pFus-faces ($n = 9$); (**b**) Average percentage signal change in object-selective ROIs; Right pFus/OTS ($n = 6$), Right LO ($n = 9$), Left pFus/OTS ($n = 7$), Left LO ($n = 7$).

comparable ratings of face-likeness as during CH sampling. The key difference here was that silhouettes in each block were defined in equal-size regions in face space such that image variability was matched across face-likeness levels (MV sampling, see Fig. 3b and Materials and Methods). As anticipated, and in stark contrast to the CH sampling results, responses in face-selective ROIs increased with the perceived face-likeness of the stimuli (Fig. 7a). That is, blocks with higher face-like ratings of silhouettes produced higher responses in face-selective regions. Average responses across nine subjects were significantly positively correlated with the average perceived face-likeness of stimuli in each of these face-selective ROIs ($r_{\text{Right mFus-faces}} = 0.83$, $P < 0.01$; $r_{\text{Right pFus-faces}} = 0.97$, $P < 10^{-5}$; $r_{\text{Right IOG-faces}} = 0.87$, $P < 0.01$; $r_{\text{Left pFus-faces}} = 0.91$, $P < 0.001$). In contrast, mean responses in object-selective ROIs were not significantly modulated by perceived face-likeness of silhouettes (all Ps > 0.1; Fig. 7b) and were

consistently high for all blocks. Thus, in MV-sampling we find a correlation between average face-likeness ratings and average responses of face-selective regions specifically. Our voxel-count measure within each ROI produced similar results: more face-selective voxels were active during high face-likeness blocks than during low face-likeness blocks, but the number of active voxels in object-selective ROIs was unaffected by perceived face-likeness (see Supporting Information Fig. 6).

A Powerful Effect of Image Variability on Face- and Object-Selective Responses

At first glance, the response patterns of face-selective ROIs during CH vs. MV sampling of face-space appear contradictory: with CH sampling, average responses across subjects in face-selective regions were negatively correlated with perceived face-likeness, whereas with MV sampling, the correlation was highly positive. We suggest that the apparent contradictory responses are due to the differences in image variability across experiments. To test this possibility and quantify the role of image variability, we examined responses across the seven subjects who participated in both studies as a function of image variability. We found that blocks with more variable silhouettes images elicited higher responses than blocks with less variable (more similar) silhouettes, across most face- and object-selective ROIs (see Fig. 8). In right-hemisphere object-selective ROIs, average responses across subjects were significantly correlated with average image variability of stimuli in a block across the 18 conditions in the CH and MV sampling experiments ($r_{\text{Right pFus/OTS}} = 0.82$, $P < 10^{-7}$; $r_{\text{Right LO}} = 0.93$, $P < P < 10^{-8}$), whereas these correlations in left-hemisphere object-selective ROIs were smaller ($r_{\text{Left pFus/OTS}} = 0.33$, $P = 0.2$; $r_{\text{LO}} = 0.54$, $P = 0.02$; see Fig. 8b). Average responses in right-hemisphere face-selective regions were also positively correlated with image variability, but to a lesser degree than in right object-selective ROIs ($r_{\text{Right mFus-faces}} = 0.76$, $P < 0.001$; $r_{\text{Right pFus-faces}} = 0.52$, $P < 0.01$; $r_{\text{Right IOG-faces}} = 0.66$, $P < 10^{-4}$; $r_{\text{Left pFus-faces}} = 0.29$, $P = 0.3$; see Fig. 8a). These data suggest that image variability accounts for some of the response modulations across the CH and MV experiments.

A Two Factor Model Explains Responses in Face-Selective Regions

We note that image variability alone cannot account for response modulations in face-selective regions, since image variability was constant across the nine MV sampling blocks and yet responses in face-selective regions were modulated. We therefore asked whether responses in face-selective ROIs across the two studies could be explained as a linear combination of two factors: (1) the mean perceived face-likeness (FL) of the stimuli and (2) the image variability (IV) in each block. We coded each of the 18 blocks across the CH and MV sampling conditions according to their average face-likeness rating and their image variability and analyzed brain responses from the 7 participants who completed both studies. In face-selective ROIs, we found that both factors contribute a significant positive weight on responses. Furthermore, image variability contributed a numerically larger weight than face-likeness (right mFus-faces: $W_{\text{IV}} = 0.24$, $W_{\text{FL}} = 0.12$; right pFus-faces: $W_{\text{IV}} = 0.20$, $W_{\text{FL}} = 0.18$; right LO-faces: $W_{\text{IV}} = 0.40$, $W_{\text{FL}} = 0.28$; left pFus-faces: $W_{\text{IV}} = 0.25$, $W_{\text{FL}} = 0.24$). Together, face-likeness and image variability explained 78% of the variance of the responses across the 18 blocks in right mFus-faces, 87% in right pFus-faces, 79% in right IOG-faces, and 51% in left pFus-faces (Fig. 9). In contrast, responses in object-selective ROIs were accounted for by the single factor of image variability, as face-likeness did not contribute additional explanation of the variance of object-selective responses (Ps > 0.1). Image variability provided a better account of responses in right than left object-selective regions. It accounted for 68 and 87% of the variance in responses in right pFus/OTS and LO, respectively and 11 and 29% for the left pFus/OTS and LO.

These analyses illustrate (1) a positive effect of face-likeness on responses of face-selective regions only and (2) a positive effect of image variability on responses of both face- and object-selective ROIs. With CH sampling, these two factors were negatively correlated and produced competing effects on the responses of face-selective regions, with the effect of image variability overriding that of face-likeness. Thus, during CH sampling responses ultimately increased with image variability despite decreasing face-likeness because the weight of image variability was numerically larger. This analysis also explains the unexpected negative relationship between face-likeness and average responses in object-selective cortex during CH sampling, suggesting that this correlation was in fact driven by differences in image variability across blocks (which were lowest for the most face-like blocks).

DISCUSSION

By manipulating stimuli in a parameterized silhouette face space, we showed that two independent factors—face-likeness and image variability—drive up responses in face-selective ventral regions. Four face-selective regions (right mFus-faces, pFus-faces, IOG-faces, and left pFus-faces) showed a monotonically graded response to the perceived face-likeness of stimuli, responding more strongly to stimuli that were more face-like. However, this response profile was only evident when image variability was kept constant across blocks via MV sampling. When blocks were defined with CH sampling, and variability was negatively correlated with face-likeness, the effect of image variability dominated and responses were highest during the high-variability blocks despite their low face-likeness. Notably, this driving effect of image variability was present across multiple face- and object-selective ventral

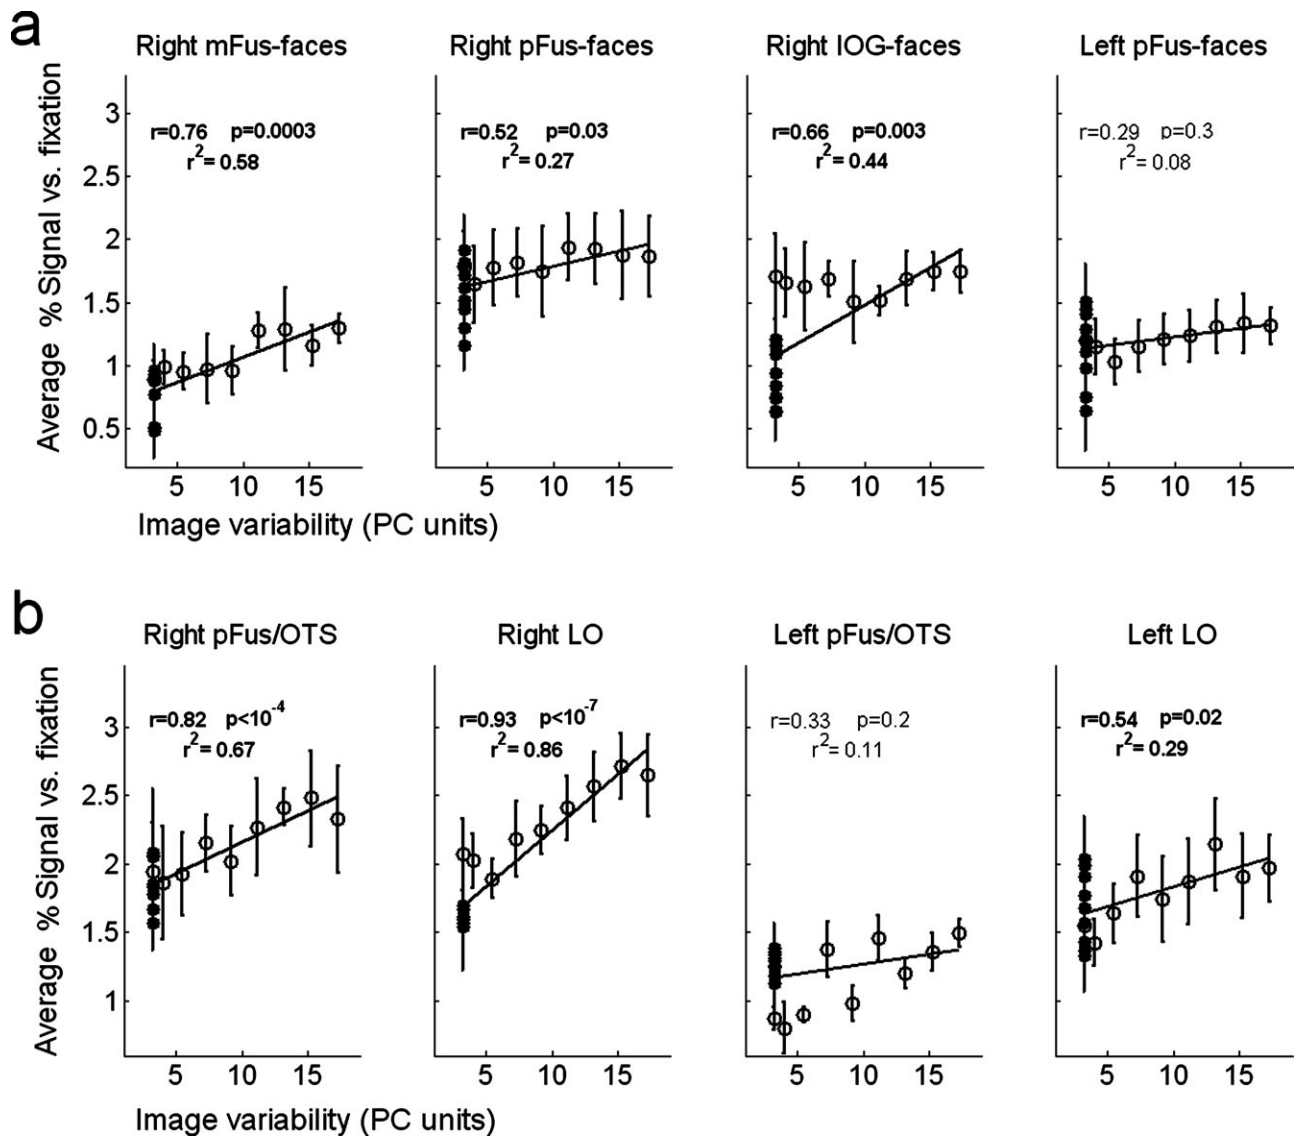

Figure 8.

Responses in face- and object-selective regions across subjects who participated in both the CH sampling (open circles) and MV sampling (black dots) experiments, plotted as a function of image variability. Each point reflects the average response across subjects to a block of stimuli. *r*- and *P*-values denote correlations and their significance, respectively, between image variability and average percentage signal across subjects; error bars denote between-subjects SEM. (**a**) Average percentage signal change across subjects in four face-selective ROIs; Number of subjects in which each ROI was localized across both studies: Right mFus-faces ($n = 6$), Right pFus-faces ($n = 7$), Right IOG-faces ($n = 4$); Left pFus-faces ($n = 7$); (**b**) Average percentage signal in object-selective ROIs; Right pFus/OTS ($n = 4$), Right LO ($n = 7$), Left pFus/OTS ($n = 3$), Left LO ($n = 5$).

regions, whereas the effect of face-likeness was only observed in face-selective ROIs.

Responses in Face-Selective Regions Track Perceived Face-Likeness

When image variability was matched across blocks, we found a striking correlation between face-likeness and average responses in face-selective regions. Moreover, these regions tracked the face-likeness of stimuli monotonically. These results have important implications for the neural mechanisms underlying face representation in the ventral stream. One possibility is that face-selective neurons in these regions show graded responses, responding more strongly to stimuli that more closely resemble a prototypical face. A second possibility is that this graded response

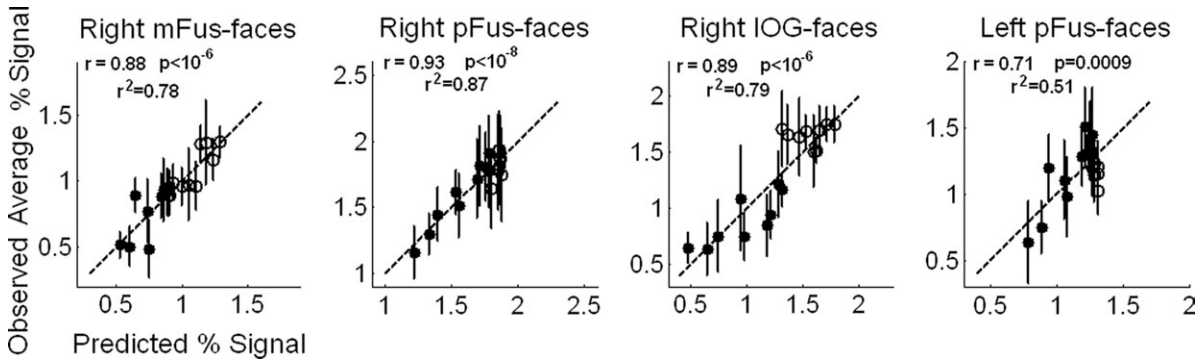

Figure 9.

Correlation between predicted responses based on the two-factor model and observed percentage signal change in face-selective ROIs across subjects who participated in both the CH sampling (open circles) and MV sampling (black dots) experiments. Each point reflects the average response across subjects to a block of stimuli. *X*-axis: predicted % signal based on two factors (face-likeness and image variability); *Y*-axis: observed average % signal across subjects; R^2 indicates proportion of variance explained across the 18 blocks from the two studies, and p-values indicate significance of the linear fit; error bars indicate between-subjects SEM; Number of subjects in which each ROI was localized across both studies: Right mFus-faces ($n = 6$), Right pFus-faces ($n = 7$), Right lOG-faces ($n = 4$); Left pFus-faces ($n = 7$).

profile is the product of different neural populations tuned to different types of stimuli [Jiang et al., 2006], with more neurons tuned to highly face-like stimuli and fewer neurons tuned to stimuli that are less face-like. These possibilities are not mutually exclusive: it may be that more neurons are tuned close to the prototypical face, and all neurons show a preference for face-like stimuli. Although our studies cannot distinguish between these possibilities, any viable neural model of face representation must account for a BOLD response that correlates with perceived face-likeness of stimuli.

Face-Specific and General Factors

Examining responses across multiple brain regions allowed us to determine which manipulations of our stimuli have face-specific consequences and which manipulations affect responses more generally across ventral occipito-temporal regions. We show that face likeness is a face-specific factor: When variability was controlled (MV sampling) face-selective responses increased with face-likeness, but object-selective responses were not modulated by face-likeness. By equating image variability across different conditions, we were able to study the effect of face-likeness in isolation and show a modulation in face-selective responses only.

In contrast, image variability is a general factor affecting responses in object- as well as face-selective regions. Responses in LO and pFus/OTS, which show no selectivity to faces, increased as a function of image variability (Fig. 8). This effect of variability was large and prevalent across ventral occipito-temporal regions. For example, in mFus-faces and pFus-faces, stimuli that were rated as least face-like elicited twice as large of a response in CH sampling than in MV sampling, because the low face-like stimuli in CH sampling were much more variable (Figs. 6 and 7, lowest face-likeness point). In fact, image variability had such a large effect on responses of face-selective regions that it reversed their preference for face-like stimuli in CH sampling, where variability was anti-correlated with face-likeness (Fig. 6).

Our data indicate that responses in face-selective regions are driven by at least two factors: face-likeness and image variability, and studying their effects requires dissociating the two factors experimentally. The face silhouette methodology is well suited for this purpose because the same shape-based parameterization can be used to manipulate a face-specific factor (in this case, face-likeness), while controlling the general factor of image variability. Face silhouettes capture a subset of the face dimensions that are relevant for face perception tasks; specifically, dimensions related to face profile shape [Davidenko, 2007]. Nevertheless, the face silhouette methodology cannot address every aspect of face representation. Previous research indicates that facial characteristics other than shape, including local features, textures, and surface reflectance information [O'Toole et al., 1999; Russell and Sinha, 2007; Schyns et al., 2002] also influence face perception. Thus, future work is needed to understand the effects of other face characteristics on face-selective responses. Critically, however, our data indicate that future examinations of the functional characteristics of face-selective regions, as well as ventral stream regions more generally, must include a method of measuring and controlling the variability of stimuli across conditions.

Effects of Different Measures of Variability in Face Silhouettes

In our stimulus space, image variability was defined as the average distance in face-space between stimuli in each

block, but other physical and perceptual measures of variability can be defined. For example, a "pixel-based" metric (see Materials and Methods) is generalizable to other stimulus spaces, as it can be defined on gray-scale or color images. For our stimuli, pixel-based variability was highly correlated with the face-space measure of image variability ($r = 0.96$, $P < 10^{-9}$ across the 18 stimulus blocks; see Supporting Information Fig. 7a) and consequently similarly explained our fMRI data. Since recent studies suggest that responses in some high-level visual regions are coupled with perceptual rather than physical variability among stimuli [Gilaie-Dotan et al., 2010; Jiang et al., 2006; Rotshtein et al., 2005], we also examined whether perceptual metrics of variability better explain our data. We considered a perceptual metric of variability by behaviorally measuring dissimilarity ratings among stimuli in each block (see Materials and Methods). We found that mean dissimilarity ratings were highly correlated with image variability ($r = 0.97$, $P < 10^{-10}$; Supporting Information Fig. 7b) as well as with pixel-based variability ($r = 0.88$, $P < 10^{-5}$; Supporting Information Fig. 7c). Thus, in our stimulus space, these three metrics of variability are tightly coupled. However, this coupling may not generalize to other stimuli. For instance, manipulating image size within a block of stimuli will result in a large pixel-based variability but relatively small perceptual variability. Further, different metrics of variability may affect various high-level visual areas differentially. Thus, in future studies, characterizing which metric of variability is relevant to control will depend on both the stimulus domain and the brain regions being investigated.

Implications of Our Data on Understanding Functional Properties of High Level Visual Cortex

Our results show that stimulus variability, whether measured in physical or perceptual units, is a critical factor to consider in fMRI investigations of high-level visual regions. Because high-level visual areas adapt to similar stimuli, responses will generally be stronger to blocks of highly variable images compared to blocks of similar images. Despite its prevalent effect in the ventral stream, many studies of face and object representation fail to consider the role of variability. When researchers examine the response properties of face-selective regions, they often focus on one factor: e.g., gender or distinctiveness. Our results indicate that focusing on one factor is insufficient unless stimuli are matched for variability. When variability is not controlled, results attributed to the manipulated factor may be misinterpreted if the factor is confounded with stimulus variability. For example, a recent study [Freeman et al., 2010] reported that FFA responds more strongly to gendered (as opposed to androgynous) blocks of faces. However, since blocks of gendered faces included both male and female faces, they were likely more variable than the blocks of androgynous faces. This raises the possibility that the apparent stronger FFA responses to gendered faces were in fact driven by image variability. Another important example is a study by Loffler et al. [2005] that used synthetic faces [Wilson et al., 2002] to manipulate face distinctiveness and found that FFA responses were stronger to blocks of distinctive faces (far from the prototype) as compared to blocks of typical faces (close to the prototype). This stronger FFA response to distinctive faces has been interpreted as evidence for a norm-based neural coding of faces in which individual neurons respond more strongly as faces deviate more from the prototype, or norm face [Panis et al., 2011; Rhodes, 2006]. However, the blocks in Loffler et al.'s study were generated using the same CH sampling method we employed in Experiment 3, rendering distinctiveness highly correlated with variability. In other words, blocks of distinctive faces were more variable than blocks of typical faces. As our data suggest, Loffler et al.'s results may have been driven by differences in variability across blocks rather than differences in face distinctiveness itself. Further, because they did not report responses in other ventral regions such as LO and pFus/OTS, it is unknown whether their effects reflect a face-specific mechanism or instead reflect a more general adaptation effect across the ventral stream. In an effort to rule out an adaptation-based explanation, Loffler and colleagues showed V1 responses were not modulated by their manipulation of distinctiveness. However, we note that early visual areas, including hV4, are not effective control regions because they do not show fMRI-A to shapes or faces [Sayres and Grill-Spector, 2006; Weiner et al., 2010]. Indeed, in the present experiment, we identified hV4 based on visual field mapping [Sayres and Grill-Spector, 2008] in five subjects in CH sampling and six subjects in MV sampling and extracted responses in each stimulus condition. First, we found no significant modulation of hV4 responses in either of the CH or MV experiments (see Supporting Information Fig. 8a,b). Second, in contrast to LO and pFus responses which show fMRI-A that depended on image variability, hV4 responses were not modulated by image variability (Supporting Information Fig. 8c). This is because image variability modulates responses only if there is fMRI-A across similar stimuli.

As researchers continue to debate the mechanisms underlying the neural coding of faces and compare the efficacy of norm-based [Leopold et al., 2006; Loffler et al., 2005; Rhodes, 2006] versus exemplar-based [Gilaie-Dotan and Malach, 2007; Gilaie-Dotan et al., 2010; Jiang et al., 2006] accounts, it is critical to decouple distinctiveness from image variability when examining its effect on face-selective responses. As our studies show, one way this may be accomplished is by generating matched-variability blocks of stimuli.

Our experiments show that stimulus variability affects responses across many ventral regions, not just face-selective ROIs. This suggests that measuring and controlling variability should become a critical tool in studies of face

perception as well as object representation more generally. Future studies may ask whether an object-selective region such as LO responds more strongly to animate versus inanimate objects, or whether face-selective regions respond preferentially to familiar versus unfamiliar faces. These questions cannot be well addressed without properly accounting for different degrees of variability across stimulus conditions. This challenge pertains not only to block-design experiments. While adaptation effects are smaller across events separated temporally relative to blocks, they are nevertheless present and persist across many intervening stimuli [Andresen et al., 2009; Sayres and Grill-Spector, 2006; Weiner et al., 2010]. Although a recent method of carry-over design [Aguirre, 2007] partially ameliorates immediate adaptation effects across stimuli, it does not control for adaptation effects across many intervening stimuli. Therefore, the effect of image variability is important to consider in any experiment in which responses are measured across a group of stimuli, including event-related experiments and multi-voxel pattern analyses. Determining appropriate measures of image variability and describing their effects on the responses of different brain regions will remain a key research goal for future neuroimaging studies of high-level visual regions.

ACKNOWLEDGMENTS

The authors thank Nathan Witthoft, Jon Winawer, and Golijeh Golarai for useful discussions during the design of the experiments and preparation of the manuscript.

REFERENCES

Aguirre J (2007): Continuous carry-over designs for fMRI. Neuroimage 35:1480–1494.

Andresen DR, Vinberg J, Grill-Spector K (2009): The representation of object viewpoint in human visual cortex. Neuroimage 45:522–536.

Andrews TJ, Ewbank MP (2004): Distinct representations for facial identity and changeable aspects of faces in the human temporal lobe. Neuroimage 23:905–913.

Andrews TJ, Schluppeck D, Homfray D, Matthews P, Blakemore C (2002): Activity in the fusiform gyrus predicts conscious perception of Rubin's vase-face illusion. Neuroimage 17:890–901.

Avidan G, Harel M, Hendler T, Ben-Bashat D, Zohary E, Malach R (2002): Contrast sensitivity in human visual areas and its relationship to object recognition. J Neurophysiol 87:3102–3116.

Barton JJS, Press DZ, Keenan JP, O'Connor M (2002): Lesions of the fusiform face area impair perception of facial configuration in prosopagnosia. Neurology 58:71–78.

Brainard DH (1997): The psychophysics toolbox. Spatial Vis 10:433–436.

Davidenko N (2007): Silhouetted face profiles: a new methodology for face perception research. J Vis 7:6.

Davidenko N (2007): The use of parameterized stimulus spaces for the study of face representation. Dissertation Abstracts International: Section B: The Sciences and Engineering.67:5434.

Davidenko N, Ramscar MJA (2006): The distinctiveness effect reverses when using well-controlled distractors. Vis Cogn 14:89–92.

Davidenko N, Witthoft N, Winawer J (2007): Gender aftereffects in face silhouettes reveal face-specific mechanisms. Vis Cogn 16:99–103.

Downing PE, Chan AW-Y, Peelen MV, Dodds CM, Kanwisher N (2006): Domain specificity in visual cortex. Cerebral Cortex 16:1453–1461.

Ewbank MP, Andrews TJ (2008): Differential sensitivity for viewpoint between familiar and unfamiliar faces in human visual cortex. Neuroimage 40:1857–1870.

Freeman JB, Rule NO, Adams RB, Ambady N (2010): The neural basis of categorical face perception: Graded representations of face gender in fusiform and orbitofrontal cortices. Cereb Cortex 20:1314–1322.

Gauthier I, Tarr MJ, Moylan J, Skudlarski P, Gore JC, Anderson AW (2000): The fusiform "face area" is part of a network that processes faces at the individual level. J Cogn Neurosci 12:495–504.

Gilaie-Dotan S, Malach R (2007): Sub-exemplar shape tuning in human face-related areas. Cereb Cortex 17:325–338.

Gilaie-Dotan S, Gelbard-Sagiv H, Malach R (2010): Perceptual shape sensitivity to upright and inverted faces is reflected in neuronal adaptation. NeuroImage 50:383–395.

Glover GH (1999): Simple analytic spiral K-space algorithm. Magn Reson Med 42:412–415.

Golarai G, Ghahremani DG, Whitfield-Gabrieli S, Reiss A, Eberhardt JL, Gabrieli JD, Grill-Spector K (2007): Differential development of high-level visual cortex correlates with category-specific recognition memory. Nat Neurosci 10: 512–522.

Golarai G, Liberman A, Yoon JM, Grill-Spector K (2010): Differential development of the ventral visual cortex extends through adolescence. Front Hum Neurosci 3:80.

Golby AJ, Gabrieli JD, Chiao JY, Eberhardt JL (2001): Differential responses in the fusiform region to same-race and other-race faces. Nat Neurosci 4:845–850.

Grill-Spector K (2003): The neural basis of object perception. Curr Opin Neurobiol 13:159–166.

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

Grill-Spector K, Kushnir T, Edelman S, Avidan G, Itzchak Y, Malach R (1999): Differential processing of objects under various viewing conditions in the human lateral occipital complex. Neuron 24:187–203.

Grill-Spector K, Knouf N, Kanwisher N (2004): The fusiform face area subserves face perception, not generic within-category identification. Nat Neurosci 7:555–562.

Grill-Spector K, Henson R, Martin A (2006): Repetition and the brain: Neural models of stimulus-specific effects. Trends Cogn Sci 10:14–23.

Hasson U, Hendler T, Ben Bashat D, Malach R (2001): Vase or face? A neural correlate of shape-selective grouping processes in the human brain. J Cogn Neurosci 13:744–753.

Jiang X, Rosen E, Zeffiro T, Vanmeter J, Blanz V, Riesenhuber M (2006): Evaluation of a shape-based model of human face discrimination using FMRI and behavioral techniques. Neuron 50:159–172.

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.

Kranz F, Ishai A (2006): Face perception is modulated by sexual preference. Curr Biol 16:63–68.

Leopold DA, Bondar IV, Giese MA (2006): Norm-based face encoding by single neurons in the monkey inferotemporal cortex. Nature 442:572–575.

Li L, Miller EK, Desimone R (1993): The representation of stimulus familiarity in anterior inferior temporal cortex. J Neurophysiol 69:1918–1929.

Loffler G, Yourganov G, Wilkinson F, Wilson HR (2005): fMRI evidence for the neural representation of faces. Nat Neurosci 8:1386–1390.

Moutoussis K, Zeki S (2002): The relationship between cortical activation and perception investigated with invisible stimuli. Proc Natl Acad Sci USA 99:9527–9532.

O'Toole AJ, Vetter T, Blanz V (1999): Three-dimensional shape and two-dimensional surface reflectance contributions to face recognition: An application of three-dimensional morphing. Vis Res 39:3145–3155.

Panis S, Wagemans J, Op de Beeck HP (2011): Dynamic norm-based encoding for unfamiliar shapes in human visual cortex. J Cogn Neurosci 23:1829–1843.

Phillips PJ, Wechsler H, Huang J, Rauss P (1998): The FERET database and evaluation procedure for face recognition algorithms. Image Vis Comput J 16:295–306.

Phillips PJ, Moon H, Rizvi SA, Rauss PJ (2000): The FERET evaluation methodology for face recognition algorithms. IEEE Trans Pattern Anal Machine Intell 22:090–1104.

Puce A, Allison T, Gore JC, McCarthy G (1995): Face-sensitive regions in human extrastriate cortex studied by functional MRI. J Neurophysiol 74:1192–1199.

Rhodes G, Jeffery L (2006): Adaptive norm-based coding of facial identity. Vis Res 46:2977–2987.

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.

Russell R, Sinha P (2007): Real world face recognition: The importance of surface reflectance properties. Perception 36:1368–1374.

Sayres R, Grill-Spector K (2006): Object-selective cortex exhibits performance-independent repetition suppression. J Neurophysiol 95:995–1007.

Sayres R, Grill-Spector K (2008): Relating retinotopic and object-selective responses in human lateral occipital cortex. J Neurophysiol 100:249–267.

Schyns PG, Bonnar L, Gosselin F (2002): Show me the features! Understanding recognition from the use of visual information. Psychol Sci 13:402–409.

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

Tong F, Nakayama K, Moscovitch M, Weinrib O, Kanwisher N (2000): Response properties of the human fusiform face area. Cogn Neuropsychol (Special Issue: The cognitive neuroscience of face processing) 17:257–279.

Tootell RBH, Devaney KJ, Young JC, Postelnicu G, Rajimehr R, Ungerleider LG (2008): fMRI mapping of a morphed continuum of 3D shapes within inferior temporal cortex. Proc Natl Acad Sci USA 105:3605–3609.

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.

Weiner KS, Grill-Spector K (2010): Sparsely-distributed organization of face and limb activations in human ventral temporal cortex. Neuroimage 52:1559–1573.

Weiner K, Sayres R, Vinberg J, Grill-Spector K (2010): fMRI-adaptation and category selectivity in human ventral temporal cortex: Regional differences across time scales. J Neurophysiol 103:3349–3365.

Wilson HR, Loffler G, Wilkinson F (2002): Synthetic faces, face cubes, and the geometry of face space. Vis Res 42:2909–2923.

Worsley KJ, Marrett S, Neelin P, Vandal AC, Friston KJ, Evans AC (1996): A unified statistical approach for determining significant signals in images of cerebral activation. Hum Brain Mapp 4:58–73.

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