

Prevalence of Selectivity for Mirror-Symmetric Views of Faces in the Ventral and Dorsal Visual Pathways

Tim C. Kietzmann,^{1,2} Jascha D. Swisher,¹ Peter König,² and Frank Tong¹

¹Psychology Department and Vanderbilt Vision Research Center, Vanderbilt University, Nashville, Tennessee 37240, and ²Institute of Cognitive Science, University of Osnabrück, 49076 Osnabrück, Germany

Although the ability to recognize faces and objects from a variety of viewpoints is crucial to our everyday behavior, the underlying cortical mechanisms are not well understood. Recently, neurons in a face-selective region of the monkey temporal cortex were reported to be selective for mirror-symmetric viewing angles of faces as they were rotated in depth (Freiwald and Tsao, 2010). This property has been suggested to constitute a key computational step in achieving full view-invariance. Here, we measured functional magnetic resonance imaging activity in nine observers as they viewed upright or inverted faces presented at five different angles (-60 , -30 , 0 , 30 , and 60°). Using multivariate pattern analysis, we show that sensitivity to viewpoint mirror symmetry is widespread in the human visual system. The effect was observed in a large band of higher order visual areas, including the occipital face area, fusiform face area, lateral occipital cortex, mid fusiform, parahippocampal place area, and extending superiorly to encompass dorsal regions V3A/B and the posterior intraparietal sulcus. In contrast, early retinotopic regions V1–hV4 failed to exhibit sensitivity to viewpoint symmetry, as their responses could be largely explained by a computational model of low-level visual similarity. Our findings suggest that selectivity for mirror-symmetric viewing angles may constitute an intermediate-level processing step shared across multiple higher order areas of the ventral and dorsal streams, setting the stage for complete viewpoint-invariant representations at subsequent levels of visual processing.

Introduction

People can recognize faces and objects across a wide variety of viewing conditions, despite changes in retinal position, size, and illumination. Changes in viewing angle represent a further challenge, as large rotations of a 3D object can drastically alter the pattern of retinal input. Although people can readily recognize objects from different viewpoints, neurophysiological studies have found that the vast majority of object-selective neurons in the monkey inferotemporal cortex exhibit viewpoint-specific rather than viewpoint-invariant tuning (Perrett et al., 1991; Logothetis et al., 1995). These findings have led to the proposal that object recognition relies on multiple view-specific representations, and that the combined input of several view-specific neurons might be a necessary precursor to obtain fully view-invariant object selectivity (Bülthoff and Edelman, 1992; Logothetis et al., 1995; Perrett et al., 1998; Ullman, 1998).

Recently however, Freiwald and Tsao (2010) reported that neurons in an intermediate region of the monkey face-processing

network exhibited the peculiar property of being selective to mirror-symmetric viewing angles of faces. For instance, neurons that responded preferentially to the view of a head rotated 60° to the left were also likely to respond to a rightward rotation of 60° , but not to intermediate near-frontal views. This pattern of viewpoint symmetry can be distinguished from previous neurophysiological reports of exclusive selectivity for a single viewpoint (Perrett et al., 1991; Logothetis et al., 1995), and has been suggested to represent a key computational step toward achieving full viewpoint invariance. However, these single-unit recordings were restricted to focal regions of interest; thus, it is presently unknown whether viewpoint symmetry is a specific property of a single region in the face-processing network or whether it might be found in other visual or category-selective areas, including regions that prefer nonface stimuli such as objects and scenes.

In the present study, we investigated whether selectivity for mirror-symmetric viewing angles might also be found in the human visual system. We monitored cortical activity using functional magnetic resonance imaging (fMRI) while subjects viewed images of upright or inverted faces, taken from five different viewpoints (Fig. 1). Using multivariate pattern analysis (Haynes and Rees, 2006; Norman et al., 2006; Tong and Pratte, 2012), we then tested whether activity patterns were more similar between mirror-symmetric viewing conditions (e.g., -60 and $+60^\circ$) than between viewing angles that lacked this relationship (e.g., -60 and 0°). If so, this would imply cortical selectivity for viewpoint symmetry similar to that recently found in the monkey (Freiwald and Tsao, 2010). However, a key difference was that our pattern analytic approach did not require a region to be selective for faces or specific facial identities to exhibit mirror-symmetric selectivity.

Received Jan. 7, 2012; revised July 2, 2012; accepted July 9, 2012.

Author contributions: T.C.K., J.D.S., P.K., and F.T. designed research; T.C.K. and J.D.S. performed research; T.C.K. analyzed data; T.C.K., J.D.S., P.K., and F.T. wrote the paper.

This work was made possible by National Science Foundation Grant BCS-0642633 and National Institutes of Health (NIH) Grant R01-EY017082 (F.T.), a Fulbright scholarship (T.C.K.), NIH National Research Service Award Fellowship F32-EY019448 (J.S.), European Research Council Grant 269716 Multisense (P.K.), and NIH P30-EY008126 Center Grant to the Vanderbilt Vision Research Center. We thank Mike Pratte and Sam Ling for valuable discussions on this project, as well as Elizabeth Counterman for her help with the data acquisition.

The authors declare no competing financial interests.

Correspondence should be addressed to Tim C. Kietzmann, Institute of Cognitive Science, University of Osnabrück, Albrechtstrasse 28, 49076 Osnabrück, Germany. E-mail: tkietzma@uni-osnabrueck.de.

DOI:10.1523/JNEUROSCI.0126-12.2012

Copyright © 2012 the authors 0270-6474/12/3211763-10\$15.00/0

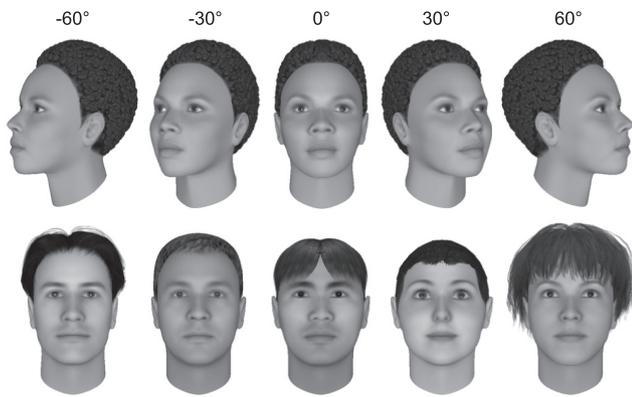


Figure 1. Stimuli. The stimuli included five different viewpoints (-60 , -30 , 0 , 30 , and 60° , upper row) of six different individuals (lower row).

To rule out potential confounding effects of low-level similarity, we developed an experimental stimulus set guided by the results of a biologically realistic model of V1 neurons (Fig. 2). We analyzed activity patterns from multiple regions of interest (ROIs) throughout the ventral and dorsal processing streams, and performed a spatially unconstrained searchlight analysis (Kriegeskorte et al., 2006) to uncover any additional areas that exhibited selectivity for viewpoint symmetry.

Materials and Methods

Participants

Ten healthy subjects (aged 22–34 years, four female) with normal or corrected-to-normal vision participated in the experiment. One subject had to be excluded from the analyses due to extreme signal dropout in the vicinity of the ear canal. All subjects were informed of their right to withdraw from the experiment at any point in time and gave written consent to participate. The study was approved by the Vanderbilt University Institutional Review Board.

Experimental design and procedure

Each experimental run consisted of 12 blocks: fixation blocks at the start and end of a run, and 10 blocks containing two presentations of each of the five viewpoint conditions (-60 , -30 , 0 , 30 , and 60°). The order of conditions was pseudorandomized and it was ensured that no condition was repeated in two consecutive blocks. Each block included three presentations of each of six face identities (pseudorandomized order) and lasted 20 s, leading to a total time of 4 min for every experimental run. Every odd run showed upright faces whereas every even run showed inverted faces. A complete scan session typically included 18 runs (9 for the upright and 9 for the inverted conditions) and lasted 2–2.5 h.

Each stimulus was shown for 800 ms, followed by a blank of 311 ms in which only a small fixation dot remained visible. Subjects were asked to perform a one-back detection task for which stimulus repetitions occurred randomly with a probability of 0.15. Furthermore, the horizontal and vertical position of the stimuli was randomly jittered by up to 10 pixels. The display computer was a luminance-calibrated MacBook Pro using MATLAB and the Psychophysics Toolbox 3 (Brainard, 1997) for experimental control. The stimuli were projected on a screen and covered 5.5° of visual angle.

Stimuli and V1 model

The stimuli were created using the face modeling software FaceGen (Singular Inversions). They included six individuals (three female) shown from five different viewpoints (-60 , -30 , 0 , 30 , and 60°) on a white background, leading to a total of 30 grayscale stimuli.

To exclude potential low-level explanations, we tested the stimuli before the experiment based on a biologically realistic model of V1 simple cells (Serre and Riesenhuber, 2004). This model is based on a set of 2D Gabor functions with 17 different receptive field sizes and four orientations, the parameters of which were previously estimated based on data

from monkey electrophysiology (Fig. 2*a*). Moreover, to account for the effects of decreasing visual acuity in the periphery of the visual field, we added an additional preprocessing step in which the stimuli were “foveated” such that the model input contained high spatial resolution only in the center of the stimulus and decreasing high spatial frequency content toward the periphery. The output of the model was used to create correlation matrices depicting the low-level similarity between the different experimental conditions given a set of stimuli. This way, we were able to evaluate whether effects of viewpoint symmetry were evident in this low-level description of the stimuli and to change them accordingly. Interestingly, an analysis of the default FaceGen stimuli, which do not include hair, revealed higher low-level similarity between the mirror-symmetric viewing angles (e.g., -60 and 60°) than between the respective angles and the straight-on face (-60 and 0°). We suspected that this low-level confound was due to the relatively homogenous and low-texture posterior part of the head. Because of this, we added structured hair to the face stimuli and were thereby able to overcome this low-level confound (Fig. 2*b,c*). In fact, the resulting faces even exhibit decreased correlations between viewpoint-symmetric viewing angles compared with the correlations with the straight-on faces. An overview of the final faces and face angles can be seen in Figure 1.

In addition to allowing us to avoid low-level confounds arising from the stimuli, we used the output of the computational V1 model for the final stimuli as a predictor of low-level similarity in the later multivariate analysis.

MRI data acquisition

The experimental data were collected at the Vanderbilt University Institute for Imaging Science using a 3 T Philips Intera Achieva MRI scanner with an eight-channel head coil. The functional data were acquired using standard gradient-echo echoplanar T2*-weighted imaging with 28 slices, aligned approximately perpendicular to the calcarine sulcus and covering the entire occipital lobe as well as the posterior parietal and posterior temporal cortex (TR, 2 s; TE, 35 ms; flip angle, 80° ; FOV, 192×192 ; slice thickness 3 mm with no gap; in-plane resolution, 3×3 mm). In addition to the functional images, we collected a T1-weighted anatomical image for every subject (1 mm isotropic voxels). A custom bite bar system was used to minimize the subject’s head motion.

fMRI analysis

Preprocessing. Preprocessing of the fMRI data was based on Freesurfer, FSL and custom MATLAB scripts. The functional data were first motion corrected with respect to the average of one run. This average image was also used to coregister the functional with the structural T1 data. After detrending the functional data and converting to percentage signal change, the spatial mean of the individual ROIs was regressed out at every point in time. Following this, we z-transformed the data with respect to the mean and SD of the signal across the whole run. Finally, to extract patterns of voxel activity for the different conditions, we took the average across the corresponding time series, excluding the first 8 s after condition onset to account for hemodynamic lag. For the ROI analyses and the searchlight estimates, no smoothing was applied and the data remained in its native space (subject coalignment was performed based on individual cortical curvature, as described below). The structural volumes were automatically segmented into gray matter and white matter and flattened/inflated using Freesurfer (Dale et al., 1999; Fischl et al., 1999a).

Correlation analysis. After preprocessing the data, we estimated the similarity of response patterns across all viewing angles in the different ROIs. To do this, we used a Pearson correlation measure along with an iterative split-half procedure. In the split-half method, individual functional runs were randomly divided into two sets and the respective average response vectors of the two halves were then used to estimate the correlations between the different conditions. The resulting correlation values were then Fisher z-transformed. This entire procedure was repeated for 2000 random splits of the functional runs for each subject. The average of the resulting correlation values can be represented as a correlation matrix, which depicts the similarity of the voxel response patterns across all pairs of conditions, including the similarity of repeated presentations of one condition with itself. The standardized correlation matrix

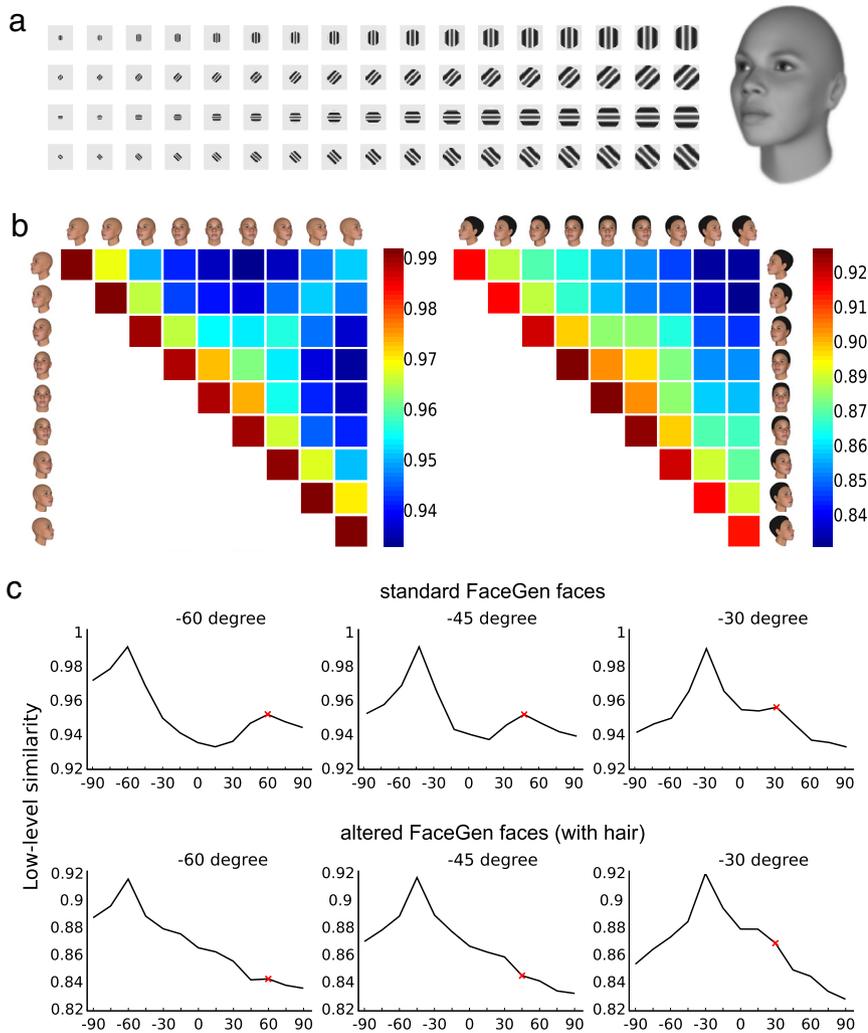


Figure 2. Control for low-level confounds. *a*, To exclude the possibility that low-level features of the stimuli would already lead to patterns of viewpoint mirror symmetry, a biologically realistic model of V1 simple cells was implemented (see Materials and Methods for details). As shown on an exemplary face on the right, the stimuli were spatially filtered (foveated) to account for differences in visual accuracy. The size of the face is proportional to the size of the Gabor filters used in the model. *b*, The V1 model responses to the standard FaceGen stimuli, as shown on the left, showed increased correlations for mirror-symmetric head orientations. This low-level confound was overcome by the addition of structured hair (shown on the right). *c*, The low-level similarity tuning curves, as estimated from the model. The red “x” marks the mirror-symmetric viewpoint.

can then be used as input for subsequent analyses, in which its congruency with different models or predictors are tested (Kriegeskorte et al., 2008a,b).

Here, we applied a two-step analysis procedure. First we estimated the extent to which the similarity of activity patterns across changes in viewpoint could be attributed to low-level similarity. Next, we estimated the degree to which viewpoint symmetry accounted for the response patterns in the region of interest, independent of that region’s sensitivity to low-level similarity.

To estimate whether an ROI showed selectivity for low-level similarity, we computed the Spearman correlation between the predicted correlation matrix of the computational V1 model (Figs. 2, 3*a*, left) and the empirical correlation matrices, which were estimated based on the activity patterns in the respective ROIs of both hemispheres of each subject. (For the analysis, only the upper triangular part of the matrices was used, as the correlation matrices are themselves essentially symmetric, with slight deviations from perfect symmetry caused by the finite number of split-half replications. Hence the lower triangular part of the matrix can be neglected in the analysis for reasons of efficiency.) This approach leads to an effect estimate for each subject and ROI. Thus, to test whether an ROI showed significant effects of low-level similarity, the corresponding

distribution of correlation values can be tested against a null result of zero correlation by applying a *t* test to the Fisher-transformed correlation values.

Next, we estimated the effects of viewpoint symmetry by constructing a model correlation matrix that predicted high correlation values for mirror-symmetric views and low correlations for nonsymmetrical ones (Fig. 3*a*, right). It should be noted that regardless of whether a brain region was selective for low-level similarity or mirror symmetry, both forms of selectivity should lead to high correlation values along the diagonal of the correlation matrix when the same viewpoint is presented. However, to be conservative in our estimates of sensitivity to mirror symmetry, we chose to use a prediction matrix with high correlations only for the mirror-symmetric cells. As a result, the model matrices for low-level similarity (Fig. 3*a*, left) and mirror symmetry (Fig. 3*a*, right) were nearly orthogonal. To ensure complete orthogonality, we first regressed out the effects of low-level similarity from the empirical matrix and then computed the effect size of viewpoint symmetry based on the residual pattern, again based on a Spearman correlation. This allowed us to calculate the partial correlation between the predicted effects of mirror symmetry and the empirical measures of cortical similarity across changes in viewpoint, having partialled out the potential contributions of low-level visual similarity (Kriegeskorte et al., 2008a).

Finally, we compared the correlation matrices of the different ROIs, independently of the two models tested. For this, we first correlated all ROI correlation matrices with each other (taking as basis the average correlation matrix across subjects and inversion condition), thereby forming a second-order similarity matrix. This matrix was then subject to a principal component analysis (PCA). Similar to multidimensional scaling, this approach allowed us to visualize the similarity relationship of the individual ROIs in a lower dimensional space. To directly estimate which parts of the correlation matrix explained most variance across ROIs, we performed an additional PCA in which we used every cell of the upper diagonal of the correlation matrix as an input dimension and the different ROIs as individual observations. This approach has the advantage that the principal components can be visualized in the same manner as the correlation matrices of the ROIs.

Searchlight analysis. The computation for the searchlight analysis was mostly identical to the correlation analysis performed for the ROIs, but based only on data from cortical voxels falling within the respective searchlight. The underlying functional data were *z*-transformed based on the mean activity and SD of each voxel across the whole run. Each searchlight included $3 \times 3 \times 3$ voxels, from which a correlation matrix was estimated using Pearson correlation and an iterative split-half approach (using 200 random splits), as described above. Once estimated, the correlation matrix was tested for its correlation with a model of low-level similarity and its partial correlation with the viewpoint symmetry model after regressing out the effects of low-level similarity. This again yielded an effect estimate for each of the two models, which were assigned to the voxel in the center of the searchlight. Shifting the searchlight across the whole brain then leads to an effect map for the two models, low-level similarity and viewpoint symmetry, for every subject. For the group analysis, the results of the searchlight analysis for all subjects were first trans-

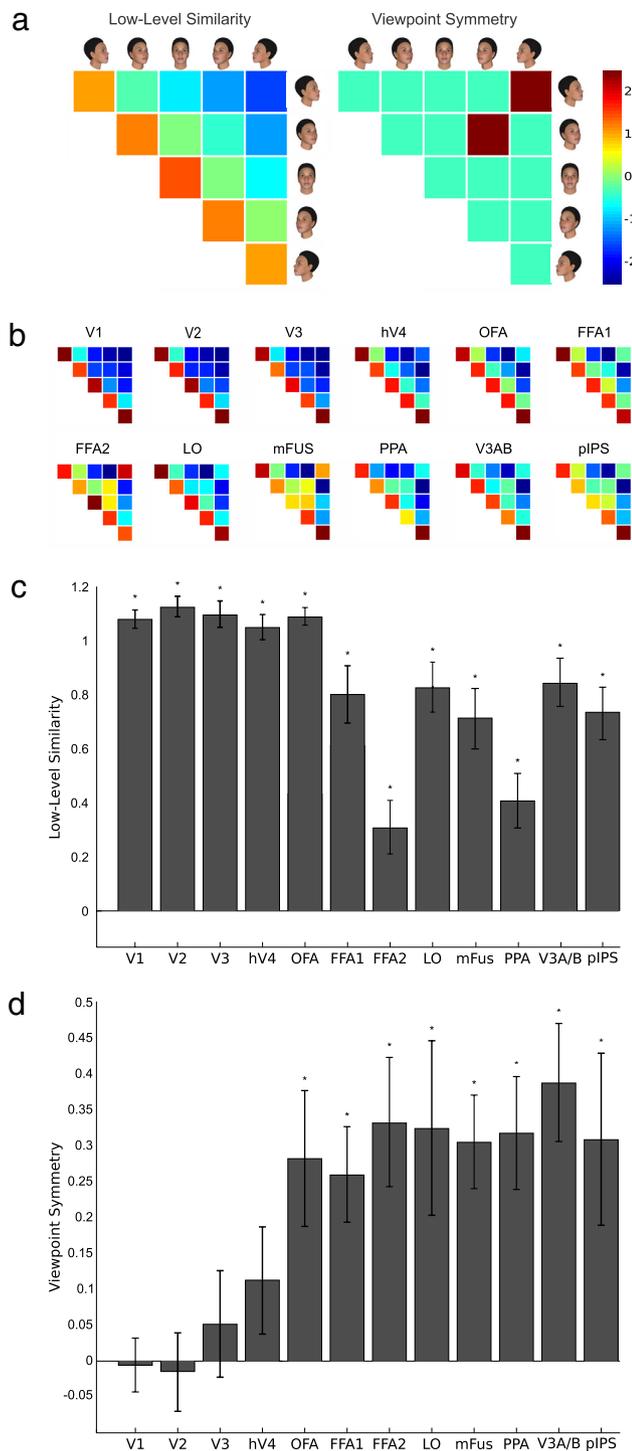


Figure 3. Effect estimates for low-level similarity and viewpoint symmetry. **a**, In a given ROI, the effects of low-level similarity were estimated by correlating the upper triangle of the empirical correlation matrices of both hemispheres with a model of low-level similarity, derived from the output of a computational V1 model (shown left). The effects of viewpoint symmetry, which predict higher correlation values for viewpoints with mirror-symmetric viewing angles, were estimated based on the partial correlation between the empirical correlation matrix and the viewpoint-symmetry model (right), after first regressing out the effects of low-level similarity. **b**, Visualization of the average correlation matrix for the ROIs. Please note that we estimated the effect sizes for every subject individually and not based on these averages. **c**, The average effect sizes of low-level similarity in the different ROIs (error bars indicate SEM). All regions show significant effects. **d**, Average effect size of viewpoint symmetry. While higher level ROIs show significant effects of viewpoint symmetry, the early and intermediate-level areas V1–hV4 do not (see text for details and p values).

formed into a common space (fsaverage) (Fischl et al., 1999b) via spherical averaging based on cortical surfaces, and smoothed to account for smaller errors due to imperfect intersubject alignment using a 6 mm full-width-half-maximum Gaussian kernel. The effects were then modeled by a general linear regression. The resulting significance map was subject to a clusterwise correction for multiple comparisons based on Monte Carlo simulations. Null volumes of normally distributed data were generated on the cortical surface and spatially filtered to match the smoothness of the subject effect size maps, as estimated by a spatial AR1 model. Clusters were defined as contiguous sets of surface vertices exceeding a significance value of $p < 0.01$. Clusters of activation were determined to be significant when their size exceeded that of the largest cluster in 95% of the simulated null volumes, for a clusterwise significance of $p < 0.05$. As an additional analysis, we performed this same procedure using a larger searchlight of $5 \times 5 \times 5$ voxels, and observed essentially the same pattern of results.

Functional ROI definitions

ROIs were defined based on independent sets of localizer data, which were either collected during the experimental session (higher order visual areas) or in a previous scan (retinotopic visual areas).

Retinotopic areas. The visual areas V1, V2, V3, hV4, V3A/B, and posterior intraparietal sulcus (pIPS) were defined based on standard retinotopic mapping (Sereno et al., 1995; Engel et al., 1997) on a flattened cortical representation. pIPS was defined as the union of areas V7 and IPS1 and 2, the borders between which could not be clearly delineated in all hemispheres. In four hemispheres (three left, one right), the retinotopic maps showed minimal significant activation in this region; for these hemispheres, approximate anatomical ROIs were defined along the medial bank of the pIPS. For one subject, no retinotopic mapping data were available. Here a V1 ROI was defined based on automated anatomical criteria (Hinds et al., 2008).

Higher order visual areas. In addition to the experimental runs, we also included three runs of a functional localizer targeting a number of higher order visual areas. The localizer included separate blocks showing objects, faces, and images of bodies (without heads) and blocks containing scrambled versions of the stimuli used in each of these three categories. As the focus of the main experiment was on the representations of different head orientations, the face localizer contained not only stimuli showing front-on faces but also the other orientations used in the main experiment (30, -30 , 60, and -60°). To avoid differences in the retinotopic extent of the scrambled and unscrambled versions of the images, we first fitted a 2D Gaussian function to the grayscale image of each stimulus. The resulting parameter estimates were then used to create a corresponding probability density function, which served as basis for the positioning of the scrambled parts of the images. The scrambled images therefore occupied approximately the same region of space as their unscrambled counterparts.

The sequence of localizer blocks was similar to the one used by Freiwald and Tsao (2010), showing a block of fixation, followed by scrambled faces, faces, scrambled objects, objects, scrambled bodies, bodies, and finally another block of fixation. The contrasts used for the individual regions are detailed below; an overview of the number of subjects for which the respective region could be defined is given in Table 1. Importantly, voxels previously labeled as belonging to one of the retinotopically organized areas (V1–hV4, V3A/B, and pIPS) were excluded from the higher order visual area definitions. For all of the higher level ROIs, we selected voxels exhibiting significantly larger activation in the respective contrast (at least $p < 0.01$, uncorrected). Similar to the early visual areas, the ROIs were defined on the flattened cortical representation of every individual subject.

In the fusiform face area (FFA) we localized the FFA voxels in the fusiform gyrus, whose activation was significantly higher for faces than objects (Kanwisher et al., 1997a). Where applicable, we assigned the labels FFA1 and FFA2 to the posterior and anterior patches of FFA, similar to (Pinsk et al., 2009). The occipital face area (OFA) was localized based on the same contrast as FFA and restricted to face-selective voxels in the occipital lobe (Puce et al., 1996; Gauthier et al., 2000). The lateral occipital cortex (LO) was defined as the set of voxels around the LO exhibiting

Table 1. Overview of the number of subjects for which the ROIs were successfully defined

V1	V2	V3	hV4	OFA	FFA1	FFA2	LO	mFus	PPA	V3A/B	piPS
lh: 9	lh:8	lh:8	lh:8	lh:6	lh:9	lh:6	lh:7	lh:9	lh:9	lh:9	lh:9
rh: 9	rh:8	rh:8	rh:8	rh:9	rh:9	rh:8	rh:9	rh:9	rh:9	rh:9	rh:9

FFA, Fusiform face area; LO, lateral occipital cortex; mFus, mid fusiform; OFA, occipital face area; piPS, posterior intraparietal sulcus.

significantly higher activation for complete objects as compared with scrambled ones (Malach et al., 1995; Kanwisher et al., 1997b).

The parahippocampal place area (PPA) responds preferably to images of houses or scenes, as compared with faces (Aguirre et al., 1998; Epstein and Kanwisher, 1998). As our localizer did not include corresponding stimulus conditions, a second set of localizer data was used to define PPA. These data were collected during a different scan session. PPA was defined as voxels around the posterior parahippocampal gyrus, showing significantly higher activation for houses than faces. Mid fusiform gyrus (mFus) was previously described as an object-selective region located on the medial side of the fusiform gyrus (Grill-Spector, 2003). In line with this definition, we determined mFus as the set of voxels in the fusiform gyrus and intermediate to FFA and PPA that exhibited significantly larger activation for objects, as compared with faces. The definition of mFus excluded voxels previously defined as PPA.

Eye-tracking analysis

During the experimental runs, the subjects were asked to remain fixated on the center of the screen, as indicated by a small red dot, and their eye position was monitored using an fMRI-compatible 60 Hz eye-tracking system (Applied Science Laboratories; Eye-Trac 6). To exclude the possibility that any residual differences in eye position could explain our observed fMRI results, the available eye-tracking data of seven of our subjects were used for further analyses. First, we tested whether there were systematic differences in the average eye position in the different conditions. For this, the average horizontal and vertical gaze direction of every subject and condition was entered into a repeated-measures ANOVA. In addition to this, we adapted a similar approach to Harrison and Tong (2009) and performed the same analysis on the eye tracking data as we did previously on the fMRI data to see whether the pattern of eye movement would lead to effects in the direction of the found fMRI results. Accordingly, we estimated empirical correlation matrices based on patterns of eye movements rather than fMRI activation for every subject, and then tested their correlation with the model of low-level similarity and the partial correlation with the viewpoint symmetry model. To estimate the empirical correlation matrices, we first converted the eye-tracking data of each subject, run, and condition into a probability density function of fixation. Based on these distributions, and analogous to the fMRI analysis, we then computed a correlation matrix for every subject across all conditions by applying an iterative split-half procedure. The resulting correlation matrices, one for every subject, were then used to estimate the effects of low-level similarity and viewpoint symmetry.

Results

Participants viewed upright or inverted faces in separate experimental runs, shown from five possible viewpoints using a randomized fMRI block design. Observers were instructed to maintain fixation on a point in the center of the screen and to perform a one-back stimulus repetition detection task (average hit rate 69%, $d' = 2.99$).

Multivariate pattern analysis

For each region of interest, we measured the similarity of cortical activity patterns across the five presented viewing angles of faces by dividing each participant's set of fMRI runs into separate halves and measuring the correlation strength (Fisher z -transformed Pearson r) across these independent datasets (see Materials and Methods). All pairwise correlations between face

viewpoints can be displayed in a correlation matrix, which can then be used as basis for further analyses. Here, we analyzed the pattern of correlations across the different viewpoints, based on two separate models of visual selectivity. The first model estimated low-level visual similarity, based on a computational V1 simple cell model (Fig. 3*a*, left). This model predicts that repeats of the same viewpoint will elicit high correlation values (i.e., similar pattern), nearby viewpoints will elicit moderate correlation values, and distal mirror-symmetric views will elicit low correlation values. In contrast, the second model, viewpoint symmetry, predicts increased correlation values between mirror-symmetric views as compared with nonsymmetrical ones (Fig. 3*a*, right). Our measure of sensitivity to mirror-symmetric viewpoints focused specifically on the predicted relationship between +60 and –60° views, and between +30 and –30° views. With this, we ensured that our measures of goodness-of-fit with this model were approximately orthogonal to our measures of sensitivity to low-level similarity. For a given ROI, the contributions of the two effects, low-level similarity and viewpoint symmetry, were assessed based on the agreement of the respective model with the empirical correlation matrices across the individual subjects (see Materials and Methods). As an overview, the average correlation matrices of the ROIs are shown in Figure 3*b*.

Using this approach, we first assessed the effects of low-level similarity and viewpoint symmetry in early visual areas V1–hV4. Our analysis of these areas revealed significant effects of low-level similarity ($p < 0.001$ in all cases, one-tailed t test), but no significant viewpoint-symmetric effects ($p = 0.55$, $p = 0.6$, $p = 0.64$, and $p = 0.08$ for V1, V2, V3, and hV4, respectively; one-tailed t test). Low-level similarity alone, as predicted by our computational V1 model, accounted for 72, 75, 72, and 65% of the total variance for V1, V2, V3, and hV4, respectively.

Following this, we analyzed higher order face-selective regions including the OFA, as well as the posterior and anterior segments of the FFA (FFA1 and 2). Again, all these regions showed significant patterns of low-level similarity ($p < 0.02$ in all cases, one-tailed t test; Fig. 3*c*). In contrast to the early visual areas, however, they also exhibited reliable effects of viewpoint symmetry ($p < 0.01$, one-tailed t test; Fig. 3*d*). Moreover, although LO, mFus, and PPA are known to respond maximally to views of objects or scenes, they are nevertheless activated by stimuli showing faces (Ishai et al., 1999). We therefore investigated whether activation patterns in these ROIs might also reveal effects of viewpoint symmetry. All three areas, LO, mFUS, and PPA, showed reliable effects of viewpoint symmetry ($p < 0.02$ in all cases), as well as low-level similarity ($p < 0.02$ in all cases). Finally, we concentrated on areas in the dorsal stream of visual processing and tested areas V3A/B and the piPS. Both regions showed significant effects of low-level similarity ($p < 0.001$) as well as viewpoint symmetry ($p < 0.02$).

To summarize, we found statistically significant effects of low-level visual similarity as well as viewpoint symmetry in all tested higher order visual areas. In contrast, early and intermediate-

level visual areas only showed significant effects of low-level visual similarity and no signs of viewpoint symmetry.

We substantiated our above conclusions with a series of controls. First, we confirmed that similar results were obtained by using a standardized regression approach in which the empirical correlation matrix was jointly predicted by the model of low-level similarity and viewpoint symmetry. This approach resulted in the same pattern of significant and nonsignificant ROIs as the two-step (partial) correlation procedure: significant effects of low-level similarity across all ROIs ($p < 0.01$ in all cases, one-tailed t test) and significant effects of viewpoint symmetry for all higher order visual areas reported above ($p < 0.05$, one-tailed t test) but no such significant effects in early areas ($p > 0.05$, one-tailed t test). Second, we find no evidence for viewpoint symmetry in our computational V1 model as well as the fMRI pattern of responses in early visual areas, indicating that our choice of stimuli effectively avoided low-level stimulus confounds (Figs. 2*b*, right, 3*d*). We observed strong effects in low-level similarity in area V1, whereas effects of viewpoint symmetry emerged only at higher stages of visual processing. Next, we tested a uniform correlation matrix with additive Gaussian noise in the same protocol as the ROIs to exclude explanations based on a fully viewpoint-invariant representation or a potential bias toward one of the two model predictors. This led to no significant similarity or viewpoint symmetry effects ($p = 0.76$ and $p = 0.71$, respectively, one-tailed t test). Furthermore, our results also cannot be explained based on overall amplitude changes in the ROIs, as we regressed out the spatial average of each ROI during preprocessing (see Materials and Methods) and because the correlation measure disregards the spatial mean of every condition. To investigate the possibility that residual eye movements of our subjects could explain our effects, we examined the eye-tracking data, for which we obtained reliable measurements from seven of our nine subjects. We found no reliable differences in the mean horizontal or vertical fixation position across conditions and subjects ($p > 0.26$ in all cases, repeated-measures ANOVA with condition as factor). In addition, we analyzed the complete set of eye-tracking data based on the same approach previously applied to the fMRI data. However, we now used probability density functions of fixation to estimate the empirical correlation matrix instead of fMRI activation patterns (see Materials and Methods). This analysis showed no significant effects of low-level similarity ($p = 0.2$, one-tailed t test) or viewpoint symmetry ($p = 0.9$, one-tailed t test) in the residual eye movements. Thus, we have no indication that eye movements can account for the observed effects.

Having found patterns of viewpoint symmetry in the responses to upright faces in both face- and object-selective higher order areas, we next asked whether similar effects could also be observed following the presentation of inverted faces. Compared with the processing of upright faces, inverted faces have been suggested to rely on distinct cognitive mechanisms by engaging object- and scene-selective regions in addition to the highly specialized face-processing network (Aguirre et al., 1999; Haxby et al., 1999; Epstein et al., 2006; Pitcher et al., 2011). Hence, it was interesting to test whether the effects of viewpoint symmetry would generalize to the processing of inverted faces. Applying the same analyses as before, we found the same pattern of results for inverted faces (Fig. 4) as for upright faces (Fig. 3, compare *c*, *d*). A repeated-measures ANOVA with ROI and inversion as within-subject factors revealed a significant effect of ROI ($p < 0.05$), but no significant effects of inversion ($p = 0.38$) or interaction effects ($p = 0.34$; all p values Greenhouse–Geisser corrected). When

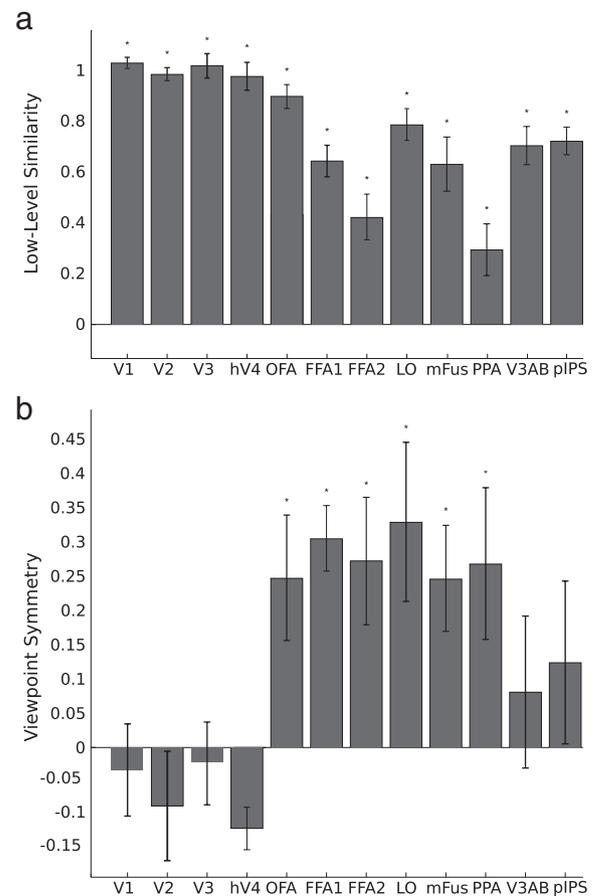


Figure 4. *a*, *b*, Effects of low-level similarity (*a*) and viewpoint symmetry (*b*) during the processing of inverted faces. As in the upright condition, all areas show significant effects of low-level similarity, whereas only higher level areas show robust effects of viewpoint symmetry.

tested individually, the higher level areas OFA, FFA1 and 2, LO, mFus, and PPA again showed significant viewpoint symmetry effects for inverted faces ($p < 0.025$ in all cases, one-tailed t test). In contrast to this, all of the early visual areas (V1–hV4) again failed to show significant effects of viewpoint symmetry ($p > 0.6$ in all cases, one-tailed t test). V3A/B and pIPS failed to show significant viewpoint symmetry effects for inverted faces ($p > 0.05$ in both cases, one-tailed t test), but there was also no statistically significant difference between the effect sizes for upright and inverted faces ($p > 0.05$ in both cases, paired t test). These results indicate that effects of mirror symmetry were also prevalent for face views presented upside-down.

To compare the correlation matrices of the different ROIs with each other, we projected an across-ROI similarity matrix into two dimensions using PCA (see Materials and Methods). In the resulting space, the distances between ROIs resemble the similarity of the respective correlation matrices (Fig. 5*a*). Notably, the first principal component, which explained 66% of the similarity structure across ROIs, exhibits a clear separation between regions with and without viewpoint symmetry (the second component explained an additional 29% of variance). Following this, we determined which cells of the correlation matrix accounted for the most variance across ROIs by performing a PCA directly on the entries of the correlation matrices (see Materials and Methods). The resulting first principal component, based on which 63% of the variance across ROIs could be explained, ex-

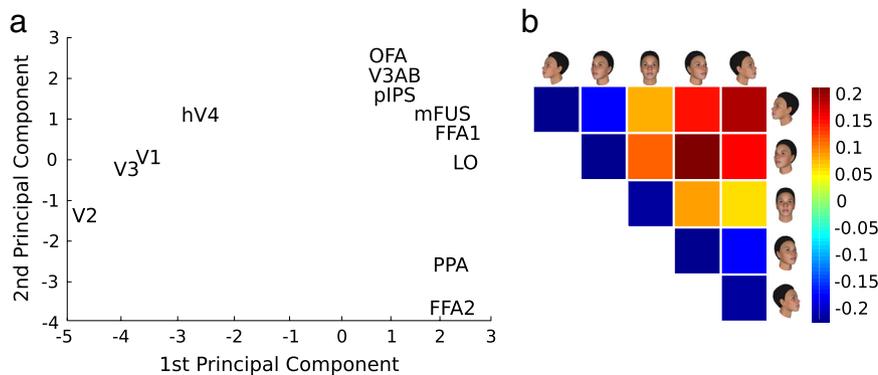


Figure 5. Principal component analyses. *a*, When projected into two dimensions, the similarity of the correlation matrices of the different ROIs can be visualized. The first component, which already explains 81.8% variance, shows a clear separation between ROIs with and without effects of viewpoint symmetry (the second component explains an additional 13.9% of the variance). *b*, The resulting first component when computing a PCA directly on the entries of the correlation matrices. The component exhibits large weights in the two diagonals, in direct agreement with the effects of low-level similarity, and viewpoint symmetry.

hibits large weights in the two matrix diagonals. This is in direct agreement with our models of low-level similarity and viewpoint symmetry (Fig. 5*b*). It should be noted that the principal component was found by simply maximizing the explained variance across ROIs; this approach is model free and did not make any assumptions regarding how the correlation matrices should tend to vary across ROIs. The results of the PCA analysis provide further confirmation that sensitivity to mirror symmetry is a prominent functional organization principle that accounts for changes in visual selectivity across the visual hierarchy.

Searchlight analysis

In addition to the analyses of specific ROIs, we performed a searchlight analysis to test for viewpoint-symmetric response patterns throughout the whole functional volume (Kriegeskorte et al., 2006). The underlying analysis was identical to the multivariate pattern analysis described above. However, instead of selecting all voxels in an ROI, the activation patterns of a local 3×3 neighborhood of voxels were used for the analysis. Shifting this searchlight across the functional volume thus yields an effect map for every subject and model. These individual subject maps were then normalized to a common space, smoothed, and tested for clusters showing significant effects of viewpoint symmetry or low-level similarity on the population level (see Materials and Methods).

The searchlight analysis revealed a large band of cortical regions exhibiting significant viewpoint-symmetric response patterns (Fig. 6). In line with our earlier results, this band of viewpoint symmetry overlapped with the previously tested ROIs, including higher order visual areas of the ventral and dorsal streams. Early retinotopic areas failed to exhibit symmetry effects and their responses were again found to be best explained based on low-level similarity. Moreover, while our searchlight analysis revealed a cluster of significant viewpoint-similarity effects in posterior superior temporal sulcus (STS), the searchlight approach revealed no significant effects of viewpoint symmetry in this region. Finally, more anterior regions such as the anterior part of the temporal lobe did not show any significant patterns of low-level similarity or viewpoint symmetry.

As a control, we performed the same analysis with a larger searchlight of $5 \times 5 \times 5$ voxels. The brain regions implicated were the same as those identified using the $3 \times 3 \times 3$ searchlight, except that the band of cortical regions was somewhat larger due

to the use of a larger searchlight. This verifies that our results were not dependent on a specific searchlight size. For this study, we present the results of the $3 \times 3 \times 3$ voxel searchlight, as this analysis provided a more conservative and spatially precise measure of the regions that displayed a preference for mirror-symmetric views of faces.

Discussion

Our analyses of cortical activity patterns revealed a spatially distributed, yet functionally specific representational property in the human visual system: selectivity for mirror-symmetric viewing angles of faces. This property was not restricted to a single focal region, but instead was found to be prevalent in a large band of higher order visual areas. In addition to regions typically associated with face processing,

OFA, FFA1, and FFA2, we observed effects of viewpoint symmetry in several cortical areas that do not respond preferentially to faces, including object-selective (LO, mFUS) and scene-selective areas (PPA). These effects were equally prevalent for inverted faces as for upright faces, even though stimulus inversion is known to impair face-specific processing (Yin, 1969; Valentine, 1988; Kanwisher et al., 1998) and the robustness of face-specific responses (Freiwald et al., 2009). This suggests that our fMRI measures of sensitivity to viewpoint symmetry do not depend on a cortical specialization for faces.

An unexpected finding was the fact that the dorsal regions V3A/B and the pIPS also exhibited selectivity for symmetric views of face stimuli. Object processing is commonly believed to rely on the ventral visual pathway (Mishkin and Ungerleider, 1982; Goodale and Milner, 1992). However, a few studies have demonstrated the presence of shape selectivity and view-invariant object selectivity in the parietal lobe as well (Serenio and Maunsell, 1998; Konen and Kastner, 2008; Króliczak et al., 2008).

Importantly, all of the visual areas we found to be sensitive to mirror-symmetric viewing angles also revealed strong effects of low-level similarity. This indicates that these ROIs failed to show complete viewpoint invariance, and that the acquisition of partial view invariance does not preclude the possibility of maintaining sensitivity to low-level image similarity as well. These findings are in line with earlier fMRI work demonstrating viewpoint-dependent adaptation effects for faces in the posterior fusiform region, including the FFA (Grill-Spector et al., 1999; Pourtois et al., 2005; Andresen et al., 2009).

In parallel to our work, a different research group has recently reported that face-selective regions including the FFA and right STS, as well as object-sensitive area LO, exhibit effects of viewpoint symmetry (Axelrod and Yovel, 2012), while no symmetry effects were found in the OFA. Here, we specifically aimed at assessing the prevalence and generality of viewpoint symmetry effects while controlling for low-level confounds typically present in standard FaceGen stimuli. Because of this, we tested upright and inverted faces across a multitude of visual areas, including retinotopically defined early visual areas V1–hV4, ventral areas OFA, FFA1 and 2, LO, mFUS, and PPA as well as dorsal areas V3 A/B, pIPS (including V7 as well as IPS1 and 2). We find positive evidence of selectivity for mirror-symmetric views not only in the FFA, OFA, and LO, but also in medial ventral temporal areas such

as mFus and the PPA, as well as dorsal visual areas in the posterior parietal cortex.

The extensive band of higher order visual areas, for which we find sensitivity to symmetric 3D viewpoints, overlaps to a considerable extent with cortical areas previously shown to prefer symmetric 2D patterns. In particular, the LO has been found to respond more strongly to symmetric dot patterns (reflected along the vertical axis) than to random dot patterns (Sasaki et al., 2005; Tyler et al., 2005). Because we observed sensitivity to symmetric views of faces rotated away from 0°, our results cannot be explained by a general preference for visually symmetric stimuli. Nevertheless, the overlap of areas raises an interesting question as to whether sensitivity to 3D viewpoint symmetry and 2D visual symmetry might reflect a shared neural mechanism.

Another related visual property is mirror reversal. Previous studies using fMRI adaptation found that ventral visual areas exhibit invariance to mirror reversals of written text, objects, and scenes (Eger et al., 2004; Dehaene et al., 2010; Dilks et al., 2011). Consistent with these neuroimaging studies, neurophysiological recordings in monkeys have shown that left–right mirror reversals lead to more similar responses in inferotemporal neurons, when compared with stimulus reversals along the vertical dimension (Rollenhagen and Olson, 2000). Although these studies found evidence of invariance to image reversal in many object-sensitive areas, consistent with the present findings, they did not directly test for selectivity to mirror-symmetric viewing angles. Because of this, it was left as an open possibility that the reported invariance to image reversals could also be explained by fully viewpoint-invariant representations. Ruling out such explanations would require the presentation of the same objects from multiple viewpoints, including views intermediate to those realized by image reversal, as was evaluated in the current work.

Together, our findings suggest that selectivity for mirror-symmetric views may constitute an intermediate-level processing step shared across multiple higher order areas of the dorsal and ventral streams. The prevalence of such representations could set the stage for realizing viewpoint-invariant representations at subsequent stages of visual processing. Indeed, Freiwald and Tsao (2010) found that viewpoint-symmetric response properties existed in a lateral region of the monkey temporal lobe (region AL), whereas neurons in a more anterior face-selective region (called AM) exhibited complete viewpoint invariance in their selectivity for different individuals (Tanaka, 1996). Interestingly, some of the neurons in AL maintained a preference for a particular facial identity across mirror-symmetric views, suggesting that such mirror-symmetric coding might serve as an important intermediate step to developing a fully view-invariant representation.

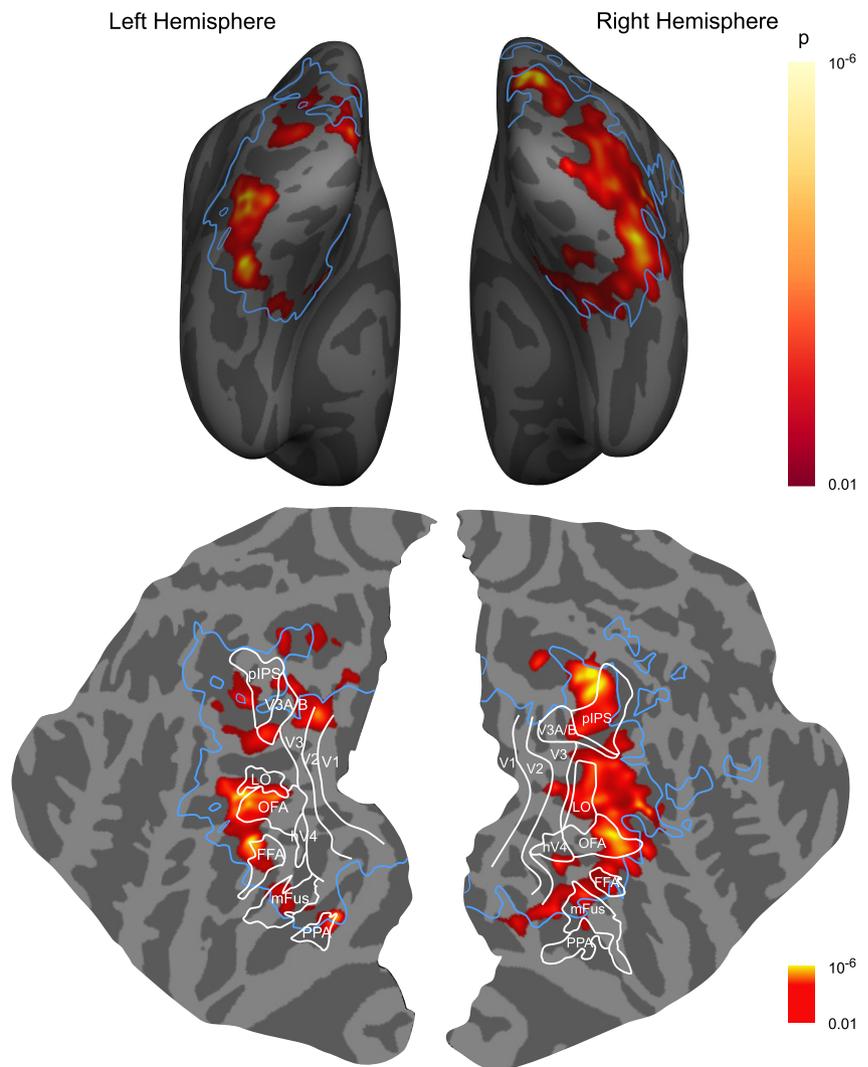


Figure 6. Searchlight results. Clusters of significant low-level similarity and viewpoint symmetry across subjects on the inflated (top) and flattened (bottom) standard brain. The light-blue line delineates regions showing significant effects of low-level similarity. Regions of significant viewpoint symmetry are marked in hot colors. They form a band of higher order visual regions, which excludes more posterior (early and intermediate-level) visual areas and more anterior areas. The delineated ROIs are from the localizer results of a representative subject (M051).

An important point to consider is why the present study found such widespread effects of viewpoint symmetry, whereas Freiwald and Tsao (2010) found these effects to be largely restricted to neurons in the anterior lateral face patch. Neurons recorded from the middle face patches (middle lateral and middle fundus of the superior temporal sulcus) often preferred a single viewpoint and failed to show evidence of viewpoint symmetry, suggesting that this visual property emerges at a relatively late stage of processing in anterior regions of the macaque visual system. Although the precise homologies between monkey and human face-selective areas have yet to be fully determined (Tsao et al., 2008), we observed effects of viewpoint symmetry at earlier processing stages in posterior ventral visual areas, including regions OFA and LO. What factors might account for the differences between studies? One major difference was that our pattern analysis approach could test for sensitivity to viewpoint symmetry without requiring the brain region in question to be selective for face stimuli or to be sensitive to facial identity. We observed strong effects of symmetry in regions such as the PPA and mFus, which prefer

objects more than faces, presumably because these areas contain partially view-invariant representations of features that are common across multiple object classes, including faces. Another factor is that fMRI pattern analysis pools information over much larger regions of cortex, and this too might have facilitated our ability to detect information about viewpoint symmetry in posterior visual areas and cortical regions that respond weakly to face stimuli. Finally, there could also be genuine differences between species, either due to innate factors or differential amounts of experience with symmetrical stimuli. Sasaki et al. (2005) tested for sensitivity to 2D symmetrical dot patterns in both humans and monkeys, and reported finding only a small region of the monkey visual cortex that showed preferential responses to symmetrical stimuli, around areas V4d, V3A, and TEO, whereas a much larger cortical region was activated in humans. Future studies might address these issues by performing comparable fMRI pattern analysis studies of viewpoint symmetry in monkeys.

It would also be interesting for future fMRI studies to investigate which regions of the human visual pathway contain view-invariant representations of facial identity. This was not possible here, as our experimental design showed different individuals from a selected viewpoint within a block. Generally, the ability to decode information about facial identity remains a major challenge for fMRI research, and although a limited degree of success has been reported (Kriegeskorte et al., 2007; Natu et al., 2010; Nestor et al., 2011) (but see Tsao et al., 2008), the present study illustrates the importance of ensuring that low-level confounds cannot account for the successful discrimination of different face stimuli.

Computationally, there could be advantages to relying on viewpoint-symmetric object representations as an intermediate processing step. View-based theories of invariant object recognition propose that viewpoint invariance can be accomplished by interpolating between a small set of informative 2D views (Poggio and Edelman, 1990; Bülthoff and Edelman, 1992; Tarr et al., 1998; Ullman, 1998; Kietzmann et al., 2009). Within such a framework, viewpoint-symmetric representations could be exploited to allow for a substantial reduction in computational complexity. While selectivity for mirror-symmetric views can itself be regarded as an example of partial viewpoint invariance, it might be particularly beneficial for encoding objects with axial symmetry, such as faces, animals, and many objects (Vetter et al., 1994). For this type of input, the number of viewpoint-specific representations required to represent an object could be substantially reduced by relying on representations that incorporate viewpoint symmetry as an intermediate processing step.

References

- Aguirre GK, Zarahn E, D'Esposito M (1998) An area within human ventral cortex sensitive to "building" stimuli: evidence and implications. *Neuron* 21:373–383.
- Aguirre GK, Singh R, D'Esposito M (1999) Stimulus inversion and the responses of face and object-sensitive cortical areas. *Neuroreport* 10:189–194.
- Andresen DR, Vinberg J, Grill-Spector K (2009) The representation of object viewpoint in human visual cortex. *Neuroimage* 45:522–536.
- Axelrod V, Yovel G (2012) Hierarchical processing of face viewpoint in human visual cortex. *J Neurosci* 32:2442–2452.
- Brainard DH (1997) The Psychophysics Toolbox. *Spat Vis* 10:433–436.
- Bülthoff HH, Edelman S (1992) Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proc Natl Acad Sci U S A* 89:60–64.
- Dale AM, Fischl B, Sereno MI (1999) Cortical surface-based analysis. I. Segmentation and surface reconstruction. *Neuroimage* 9:179–194.
- Dehaene S, Nakamura K, Jobert A, Kuroki C, Ogawa S, Cohen L (2010) Why do children make mirror errors in reading? Neural correlates of mirror invariance in the visual word form area. *Neuroimage* 49:1837–1848.
- Dilks DD, Julian JB, Kubilius J, Spelke ES, Kanwisher N (2011) Mirror-image sensitivity and invariance in object and scene processing pathways. *J Neurosci* 31:11305–11312.
- Eger E, Henson RN, Driver J, Dolan RJ (2004) BOLD repetition decreases in object-responsive ventral visual areas depend on spatial attention. *J Neurophysiol* 92:1241–1247.
- Engel SA, Glover GH, Wandell BA (1997) Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex* 7:181–192.
- Epstein RA, Higgins JS, Parker W, Aguirre GK, Cooperman S (2006) Cortical correlates of face and scene inversion: a comparison. *Neuropsychologia* 44:1145–1158.
- Epstein R, Kanwisher N (1998) A cortical representation of the local visual environment. *Nature* 392:598–601.
- Fischl B, Sereno MI, Dale AM (1999a) Cortical surface-based analysis. II: inflation, flattening, and a surface-based coordinate system. *Neuroimage* 9:195–207.
- Fischl B, Sereno MI, Tootell RB, Dale AM (1999b) High-resolution inter-subject averaging and a coordinate system for the cortical surface. *Human brain mapping* 8:272–284.
- Freiwald WA, Tsao DY (2010) Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science* 330:845–851.
- Freiwald WA, Tsao DY, Livingstone MS (2009) A face feature space in the macaque temporal lobe. *Nat Neurosci* 12:1187–1196.
- 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.
- Goodale MA, Milner AD (1992) Separate visual pathways for perception and action. *Trends Neurosci* 15:20–25.
- Grill-Spector K (2003) The neural basis of object perception. *Curr Opin Neurobiol* 13:159–166.
- Grill-Spector K, Kushnir T, Edelman S, Avidan G, Itzhak Y, Malach R (1999) Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron* 24:187–203.
- Harrison SA, Tong F (2009) Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458:632–635.
- Haxby JV, Ungerleider LG, Clark VP, Schouten JL, Hoffman EA, Martin A (1999) The effect of face inversion on activity in human neural systems for face and object perception. *Neuron* 22:189–199.
- Haynes JD, Rees G (2006) Decoding mental states from brain activity in humans. *Nat Rev Neurosci* 7:523–534.
- Hinds OP, Rajendran N, Polimeni JR, Augustinack JC, Wiggins G, Wald LL, Diana Rosas H, Potthast A, Schwartz EL, Fischl B (2008) Accurate prediction of V1 location from cortical folds in a surface coordinate system. *Neuroimage* 39:1585–1599.
- Ishai A, Ungerleider LG, Martin A, Schouten JL, Haxby JV (1999) Distributed representation of objects in the human ventral visual pathway. *Proc Natl Acad Sci U S A* 96:9379–9384.
- Kanwisher N, McDermott J, Chun MM (1997a) The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J Neurosci* 17:4302–4311.
- Kanwisher N, Woods RP, Iacoboni M, Mazziotta JC (1997b) A locus in human extrastriate cortex for visual shape analysis. *J Cogn Neurosci* 9:133–142.
- Kanwisher N, Tong F, Nakayama K (1998) The effects of face inversion on the human fusiform face area. *Cognition* 68:B1–B11.
- Kietzmann TC, Lange S, Riedmiller M (2009) Computational object recognition: a biologically motivated approach. *Biological cybernetics* 100: 59–79.
- Konen CS, Kastner S (2008) Two hierarchically organized neural systems for object information in human visual cortex. *Nat Neurosci* 11:224–231.
- Kriegeskorte N, Goebel R, Bandettini P (2006) Information-based functional brain mapping. *Proc Natl Acad Sci U S A* 103:3863–3868.
- Kriegeskorte N, Formisano E, Sorger 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.
- Kriegeskorte N, Mur M, Bandettini P (2008a) Representational similarity

- analysis - connecting the branches of systems neuroscience. *Front Syst Neurosci* 2:4.
- Kriegeskorte N, Mur M, Ruff DA, Kiani R, Bodurka J, Esteky H, Tanaka K, Bandettini PA (2008b) Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* 60:1126–1141.
- Króliczak G, McAdam TD, Quinlan DJ, Culham JC (2008) The human dorsal stream adapts to real actions and 3D shape processing: a functional magnetic resonance imaging study. *J Neurophysiol* 100:2627–2639.
- Logothetis NK, Pauls J, Poggio T (1995) Shape representation in the inferior temporal cortex of monkeys. *Curr Biol* 5:552–563.
- Malach R, Reppas JB, Benson RR, Kwong KK, Jiang H, Kennedy WA, Ledden PJ, Brady TJ, Rosen BR, Tootell RB (1995) Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proc Natl Acad Sci U S A* 92:8135–8139.
- Mishkin M, Ungerleider LG (1982) Contribution of striate inputs to the visuospatial functions of parieto-preoccipital cortex in monkeys. *Behav Brain Res* 6:57–77.
- 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.
- 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.
- Norman KA, Polyn SM, Detre GJ, Haxby JV (2006) Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10:424–430.
- Perrett DI, Oram MW, Harries MH, Bevan R, Hietanen JK, Benson PJ, Thomas S (1991) Viewer-centered and object-centered coding of heads in the macaque temporal cortex. *Exp Brain Res* 86:159–173.
- Perrett DI, Oram MW, Ashbridge E (1998) Evidence accumulation in cell populations responsive to faces: an account of generalisation of recognition without mental transformations. *Cognition* 67:111–145.
- Pinsk MA, Arcaro M, Weiner KS, Kalkus JF, Inati SJ, Gross CG, Kastner S (2009) Neural representations of faces and body parts in macaque and human cortex: a comparative fMRI study. *J Neurophysiol* 101:2581–2600.
- Pitcher D, Walsh V, Duchaine B (2011) The role of the occipital face area in the cortical face perception network. *Exp Brain Res* 209:481–493.
- Poggio T, Edelman S (1990) A network that learns to recognize three-dimensional objects. *Nature* 343:263–266.
- 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.
- Puce A, Allison T, Asgari M, Gore JC, McCarthy G (1996) Differential sensitivity of human visual cortex to faces, letterstrings, and textures: a functional magnetic resonance imaging study. *J Neurosci* 16:5205–5215.
- Rollenhagen JE, Olson CR (2000) Mirror-image confusion in single neurons of the macaque inferotemporal cortex. *Science* 287:1506–1508.
- Sasaki Y, Vanduffel W, Knutsen T, Tyler C, Tootell R (2005) Symmetry activates extrastriate visual cortex in human and nonhuman primates. *Proc Natl Acad Sci U S A* 102:3159–3163.
- Sereno AB, Maunsell JH (1998) Shape selectivity in primate lateral intraparietal cortex. *Nature* 395:500–503.
- Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB (1995) Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268:889–893.
- Serre T, Riesenhuber M (2004) Realistic modeling of simple and complex cell tuning in the HMAX model, and implications for invariant object recognition in cortex. Realistic modeling of simple and complex cell tuning in the HMAX model, and implications for invariant object recognition in cortex (CBCL Paper 239/AI Memo 2004–107). Massachusetts Institute of Technology, Cambridge, MA, August.
- Tanaka K (1996) Inferotemporal cortex and object vision. *Annu Rev Neurosci* 19:109–139.
- Tarr MJ, Williams P, Hayward WG, Gauthier I (1998) Three-dimensional object recognition is viewpoint dependent. *Nat Neurosci* 1:275–277.
- Tong F, Pratte MS (2012) Decoding patterns of human brain activity. *Annu Rev Psychol* 63:483–509.
- Tsao DY, Moeller S, Freiwald WA (2008) Comparing face patch systems in macaques and humans. *Proc Natl Acad Sci U S A* 105:19514–19519.
- Tyler CW, Baseler HA, Kontsevich LL, Likova LT, Wade AR, Wandell BA (2005) Predominantly extra-retinotopic cortical response to pattern symmetry. *Neuroimage* 24:306–314.
- Ullman S (1998) Three-dimensional object recognition based on the combination of views. *Cognition* 67:21–44.
- Valentine T (1988) Upside-down faces: A review of the effect of inversion upon face recognition. *Br J Psychol* 79:471–491.
- Vetter T, Poggio T, Bülthoff HH (1994) The importance of symmetry and virtual views in three-dimensional object recognition. *Curr Biol* 4:18–23.
- Yin R (1969) Looking at upside-down faces. *J Exp Psychol* 81:141–145.