

Dissociable Neural Patterns of Facial Identity across Changes in Viewpoint

Vaidehi S. Natu¹, Fang Jiang², Abhijit Narvekar¹, Shaiyan Keshvari³,
Volker Blanz⁴, and Alice J. O’Toole¹

Abstract

■ We examined the neural response patterns for facial identity independent of viewpoint and for viewpoint independent of identity. Neural activation patterns for identity and viewpoint were collected in an fMRI experiment. Faces appeared in identity-constant blocks, with variable viewpoint, and in viewpoint-constant blocks, with variable identity. Pattern-based classifiers were used to discriminate neural response patterns for all possible pairs of identities and viewpoints. To increase the likelihood of detecting distinct neural activation patterns for identity, we tested maximally dissimilar “face”–“antiface” pairs and normal face pairs. Neural response patterns for four of six identity pairs, including the “face”–“antiface” pairs, were discriminated at levels above chance. A behavioral experiment showed accord between percep-

tual and neural discrimination, indicating that the classifier tapped a high-level visual identity code. Neural activity patterns across a broad span of ventral temporal (VT) cortex, including fusiform gyrus and lateral occipital areas (LOC), were required for identity discrimination. For viewpoint, five of six viewpoint pairs were discriminated neurally. Viewpoint discrimination was most accurate with a broad span of VT cortex, but the neural and perceptual discrimination patterns differed. Less accurate discrimination of viewpoint, more consistent with human perception, was found in right posterior superior temporal sulcus, suggesting redundant viewpoint codes optimized for different functions. This study provides the first evidence that it is possible to dissociate neural activation patterns for identity and viewpoint independently. ■

INTRODUCTION

The human face provides information about the identity of a person and about a host of socially relevant cues to a person’s internal state and social intent. To identify a face, we must encode the information that makes it unique, or different, from all other faces. This code must generalize across two-dimensional affine transformations (e.g., size and position) and across transformations that depend on the three-dimensional shape of a face (e.g., viewpoint/pose and illumination). From the perspective of determining identity, head orientation is a “nuisance variable” that makes the task of face identification more challenging for the neural processing system. Ultimately, the neural system must solve the problem of mapping multiple, dissimilar images onto a code for an individual’s unique identity. By contrast, the social information conveyed by a face is carried in nonrigid deformations of facial shape (e.g., expressions) and in rigid changes to head orientation and gaze. These cues provide information about a person’s mood and/or their current focus of attention (cf. Haxby, Hoffman, & Gobbini, 2000), but are generally irrelevant for identifying a face.

The distributed systems framework proposed by Haxby et al. (2000) separates the processing of the invariant features useful for identifying faces from the changeable aspects of faces useful for social signaling into two neural streams. They propose that the invariant properties of faces are processed in lateral fusiform gyrus and the changeable aspects of faces are processed in posterior superior temporal sulcus (pSTS). In the present study, we consider the neural processing of identity and viewpoint. The viewpoint from which we see a face strongly constrains the visual information we can access about the face. It is likely that viewpoint is processed both in the invariant and changeable streams, albeit with different goals (cf. Fang, Murray, & He, 2006). For the task of identification, the goal of the neural processing should be to normalize or discount viewpoint in order to map disparate images onto a common code. For social signaling, the goal of the neural processing should be an accurate determination of where a person is looking. These two functions are distinct and may be supported by different types of coding mechanisms (Haxby et al., 2000).

Much is known about the neural regions important for coding face identity information from studies using fMR-adaptation (fMR-A) (cf. Ewbank & Andrews, 2008; Andrews & Ewbank, 2004; Grill-Spector et al., 1999) and repetition suppression methods (Pourtois, Schwartz, Spiridon,

¹The University of Texas at Dallas, Richardson, TX, ²Université Catholique de Louvain, Belgium, ³Rice University, Houston, TX, ⁴The University of Siegen, Germany

Martuzzi, & Vuilleumier, 2009; Eger, Schweinberger, Dolan, & Henson, 2005; Pourtois, Schwartz, Seghier, Lazeyras, & Vuilleumier, 2005a, 2005b; Rotshtein, Henson, Treves, Driver, & Dolan, 2005). In both fMR-A and repetition suppression methods, evidence for the neural coding of identity is signaled by adaptation or response suppression to repeated presentations of the same face identity. In these studies, there are three critical parameters. The first is the variability of “same identity” images with respect to viewpoint and image characteristics. “Same” identity has been defined with identical images (Henson, Shallice, & Dolan, 2000), moderately different images (Eger et al., 2005), images changed in viewpoint by small amount (2°–8°; Ewbank & Andrews, 2008), images changed substantially in viewpoint (0°–45°; Pourtois et al., 2009, 2005a), and frontal view images that are changed with morphing methods that selectively alter the physical and labeling components of face identity (cf. Gilaie-Dotan & Malach, 2007; Rotshtein et al., 2005).

A second critical parameter is the familiarity of the faces. Identity adaptation has been tested both with unfamiliar faces (Pourtois et al., 2009, 2005a; Andrews & Ewbank, 2004) and with familiar faces operationally defined as famous faces (Pourtois et al., 2005b; Rotshtein et al., 2005; Henson et al., 2000). It is widely recognized that human face perception is more robust to image and view changes for familiar versus unfamiliar faces (Hancock, Bruce, & Burton, 2000). Moreover, in addition to the increased perceptual flexibility humans show for familiar faces, there is an important difference between visually familiar faces and “famous faces.” Famous or personally familiar faces are likely to have neural codes that include visual, semantic, and emotive components.

A third parameter of these identity studies is the time course of the adaptation. Fang, Murray, and He (2006) found different patterns of release from fMR-A with longer versus shorter adaptation times. In particular, they compared longer adaptation times, typical of perceptual adaptation studies, with shorter adaptation times, typical of neuroimaging adaptation studies. They found differences in the pattern of release from face identity adaptation as a function of the degree of viewpoint change.¹

A sketch of the results from fMR-A and repetition suppression studies of face identity appears in Table 1. Although comparisons across studies are complicated due to method and stimulus differences, some common themes emerge. First, brain areas in fusiform gyrus adapt to the identity of both familiar and unfamiliar faces (see Table 1). The sensitivity of the fusiform area to variations in face identity is consistent with previous work, including findings that suggest the functionally defined fusiform face area (FFA) (Kanwisher, McDermott, & Chun, 1997) as the primary lesion site in prosopagnosia (Barton, Press, Keenan, & O’Connor, 2002; Hadjikhani & De Gelder, 2002; Damasio, Damasio, & Van Hoesen, 1982). It is further consistent with the modulation of FFA response with behavioral identification performance (Grill-Spector, Knouf, & Kanwisher,

2004) and with the preference of FFA for upright faces over inverted faces (Yovel & Kanwisher, 2005).

A second theme to emerge from fMR-A and repetition suppression studies is that identity sensitivity in FFA for unfamiliar faces is viewpoint-dependent (Pourtois et al., 2009, 2005a, 2005b; Ewbank & Andrews, 2008; Andrews & Ewbank, 2004; Grill-Spector et al., 1999) (see Table 1). It is worth noting, however, that other non-face-selective areas in the fusiform adapt to the identity of unfamiliar faces over viewpoint change (cf. the medial fusiform in the left, but not right, hemisphere, Pourtois et al., 2005a; the right medial fusiform and non-face-selective areas of the lateral occipital complex [LOC] bilaterally, Pourtois et al., 2009).

There is also evidence for some degree of image dependency for unfamiliar faces in FFA (Gilaie-Dotan & Malach, 2007; Eger et al., 2005). Thus, although the FFA response to two-dimensional affine transformations (e.g., size change) is invariant (Eger et al., 2005; Grill-Spector et al., 1999), responses to picture and morph-based face changes are not. For example, Eger et al. (2005) found no repetition suppression for identity when different images were used. Gilaie-Dotan and Malach (2007) found subcategorical identity sensitivity (i.e., failure to adapt) for view-constant face images altered by morphing methods. Specifically, FFA recovered completely from adaptation for unfamiliar faces with subtle morph-based changes that did not alter the perceived identity of the face.

For familiar (famous) faces, the adaptation and repetition suppression data are less convergent. Pourtois et al. (2005b) found viewpoint-insensitive identity processing for famous faces in left middle temporal and left inferior frontal cortex, but not in lateral fusiform gyrus. Ewbank and Andrews (2008), however, found viewpoint-invariant adaptation in FFA for famous faces. Eger et al. (2005) found greater generalization over different images of famous faces in anterior fusiform than in mid-fusiform. Differences in the viewpoint variations tested in these studies may account for the divergent results. Specifically, Pourtois et al. (2005b) tested large viewpoint changes (up to about 45°), Ewbank and Andrews tested smaller viewpoint changes (up to 8°), and Eger et al. did not control explicitly for view change. In a further study using morphed faces that did not vary in viewpoint, Rotshtein et al. (2005) found FFA sensitivity for a change in perceived identity for famous faces, but not for an equivalent amount of physical change that did not alter perceived identity.

In combination, studies with unfamiliar faces suggest that the identity information coded in face-selective areas in fusiform gyrus has limited ability to generalize across changes in viewpoint. Studies with familiar (famous) faces are currently too divergent in methods, stimuli, and results to draw firm conclusions. Although it is clear that neural code for face identity must be highly sensitive to subtle changes in faces, it must also be capable of setting boundaries around identity categories that are sufficiently tolerant to image and view variations to be useful for the task

Table 1. Face Identity Findings for fMR-A and Response Suppression Studies

<i>Familiarity</i>	<i>Viewpoint or Image</i>	<i>Brain Region</i>	<i>Source</i>
Unfamiliar faces	viewpoint dependent	FFA	Grill-Spector et al. (1999)
		fusiform gyrus	Andrews and Ewbank (2004)
		fusiform gyrus	Ewbank and Andrews (2008)
		right medial fusiform gyrus, rFFA	Pourtois et al. (2005a)
		face-selective areas (long adaptation), right fusiform area, lateral occipital complex (short-term adaptation)	Fang et al. (2006)
	viewpoint independent	right medial fusiform gyrus	Pourtois et al. (2005b)
			Pourtois et al. (2009)
		left medial fusiform gyrus	Pourtois et al. (2005a)
	image dependent	bilateral mid-fusiform, anterior fusiform (right > left)	Eger et al. (2005)
		FFA	Gilaie-Dotan and Malach (2007)
Familiar faces	viewpoint dependent	lateral fusiform cortex	Pourtois et al. (2005b)
	viewpoint independent	FFA (up to 8° of rotation)	Ewbank and Andrews (2008)
		left middle temporal, left inferior frontal cortex	Pourtois et al. (2005b)
	image dependent	bilateral fusiform gyrus	Eger et al. (2005)
	image independent	left anterior fusiform gyrus	Eger et al. (2005)
		right fusiform gyrus	Rotshtein et al. (2005)

of face recognition. Within limits, humans show flexible face perception even for unfamiliar faces, with virtually no cost for view changes less than about 15° (Valentin, Abdi, & O'Toole, 1994) and with tolerance for view change falling off gradually up to approximately 30° for a perceptual identity match (Troje & Bühlhoff, 1996). To date, most work investigating identity codes in cortex has focused on functionally defined face-selective areas, primarily in FFA, in occipital face area, or in pSTS. The results of these studies do not offer strong evidence that functionally defined face-selective areas, by themselves, can support the kinds of flexible face recognition humans show.

One difficulty with adaptation and repetition suppression methods is that they measure neural signal within predefined ROI areas independently. This limits their ability to assess interactions among areas that might collaborate to increase the robustness of identity codes. A more direct approach to exploring the neural codes for identity is to apply a pattern-based classification analysis to the task of discriminating faces by identity across a broader area of cortex. Previous studies have demonstrated that pattern classifiers can discriminate the neural codes underlying face and object categories (O'Toole, Jiang, Abdi, & Haxby, 2005; Hanson, Matsuka, & Haxby, 2004; Carlson, Schrater, & He, 2003; Cox & Savoy, 2003; Spiridon & Kanwisher, 2002; Haxby et al., 2001). More recent work indicates that these classifiers can discriminate within category object exemplars (Eger, Ashburner, Haynes, Dolan, &

Rees, 2008) and can discriminate the neural signals for two individual faces (Kriegeskorte, Formisano, Sorger, & Goebel, 2007). In the Kriegeskorte et al. (2007) study, a pattern classifier was applied to the task of discriminating the neural responses for two face images both viewed from a 45° angle. Using a searchlight model to select voxels, they were able to dissociate neural response patterns for the two faces using voxels in anterior inferotemporal (aIT) cortex, but not in FFA. Kriegeskorte et al. suggested that face detection may occur in FFA, but that individuation may engage aIT.

In the present study, we wished to investigate face identity codes that generalize across substantial changes in viewpoint. As noted, perceptual generalization across viewpoint becomes increasingly robust with familiarity. To control for the additional semantic and emotive associations that are likely to be part of the neural codes for famous faces, we tested with visually pre-familiarized faces (across viewpoint) rather than with famous faces.² We demonstrated with a behavioral experiment that people are highly accurate at matching the identity of these faces over the viewpoint changes tested in the fMRI experiment. Given the past literature indicating the limited generalizability of identity codes in face-selective areas such as fusiform, we considered a broader area of temporal cortex than has been considered in previous studies (cf. Dricot, Sorger, Schiltz, Goebel, & Rossion, 2008). This enabled us to measure codes that may reside in the interactions of

neural responses among the brain areas, including face- and object-selective areas, and offers an advantage over adaptation paradigms that are susceptible to signal dilution with increases in the size of the defined ROIs.

As a part of this study, we also considered the neural coding for viewpoint, independent of identity. As noted, viewpoint can signal a person's focus of attention, but strongly constrains the visual information available for identity. The processing of viewpoint, for its own sake, is likely to be part of the distributed neural stream dedicated to processing changeable aspects of faces (Haxby et al., 2000). We note that other changeable aspects of faces such as expression are also likely to be processed in the changeable stream. We focused on viewpoint because it is primarily a visual attribute of faces, and would not (like expression) evoke emotional associations and responses. Furthermore, evidence suggests asymmetric dependencies in the processing of identity and expression (cf. Fox, Moon, Iaria, & Barton, 2009; Ganel & Goshen-Gottstein, 2004) that are unlikely to apply to viewpoint.

From a psychological perspective, although much is known about how face recognition accuracy varies with viewpoint change, remarkably little is known about the accuracy of viewpoint perception, per se (i.e., How accurate are we at determining where a person is looking?). It is worth noting a priori that unlike the subtle codes that may be a prerequisite for face identification, from an evolutionary perspective, a coarse coding of viewpoint may be more than adequate to serve most social needs (i.e., making a rough approximation of where to direct your attention next based on where someone else is looking). In fact, most functional neuroimaging studies that manipulate facial viewpoint do so as a test of robustness for identity codes. An exception to this is a study by Pageler et al. (2003), who examined the brain areas responsive to interactions of head orientation and gaze direction. They found greater activation for frontal versus averted head and eye gaze in both fusiform gyrus and pSTS. Fusiform preferred forward gaze for all head orientations, whereas pSTS showed no preference.

In the present study, we applied pattern-based classification analyses to discriminate the fine-grained neural response patterns for facial identity independent of viewpoint and viewpoint independent of facial identity. We carried out an fMRI experiment to collect the neural activation patterns elicited in response to viewing different face identities and viewpoints. In this study, we rely strongly on stimulus-based predictions of neural discriminability for the identity and viewpoint pairs tested. These predictions were verified explicitly in two perceptual experiments conducted on different participants outside the scanner. For identity, to increase the likelihood of detecting distinct patterns of neural activation for different face identities and to generate predictions about face pair discriminability, we used stimuli generated with three-dimensional morphing software (Blanz & Vetter, 1999). We used two highly dissimilar "opposite" face pairs and four "other" face pairs that con-

sisted of normal (unaltered) faces and pairings between normal and opposite faces. We hypothesized that the neural discriminability of the two highly dissimilar face pairs would be greater than the discriminability of the other face pairs. For viewpoint discrimination, we predicted that larger viewpoint changes would be more neurally discriminable than smaller viewpoint changes.

METHODS

Stimuli

The stimulus set used for the fMRI experiment consisted of the four male faces viewed from four viewpoints ranging from the frontal (0°) to the profile (90°) in increments of 30° (i.e., 0°, 30°, 60°, and 90°; Figure 1A). The face stimuli were generated from laser scan data that included a three-dimensional shape and overlying two-dimensional reflectance map (Vetter & Troje, 1997). Two of the faces were "original" (i.e., unaltered) faces (Rows 1 and 2). The two additional identities (Rows 3 and 4) were created synthetically to be opposites of the originals, using three-dimensional morphing software developed by Blanz and Vetter (1999). In this graphic model, faces are represented in a face space that directly codes their deviation in shape and reflectance from an average face ($n = 200$). An individual face is coded as a vector in high-dimensional space originating at the average. An *antiface* or opposite face is created by morphing the original face back in the direction of the average, continuing through the average to a position equidistant on the other side of the mean (cf. Leopold, O'Toole, Vetter, & Blanz, 2001). This systematically inverts the feature values on all of the axes in the face space. The antiface appears to have features opposite to that of the original face and so is highly dissimilar to the original face from which it is created.

Images of each face from the four viewpoints were created by three-dimensional graphic rendering of the head models.

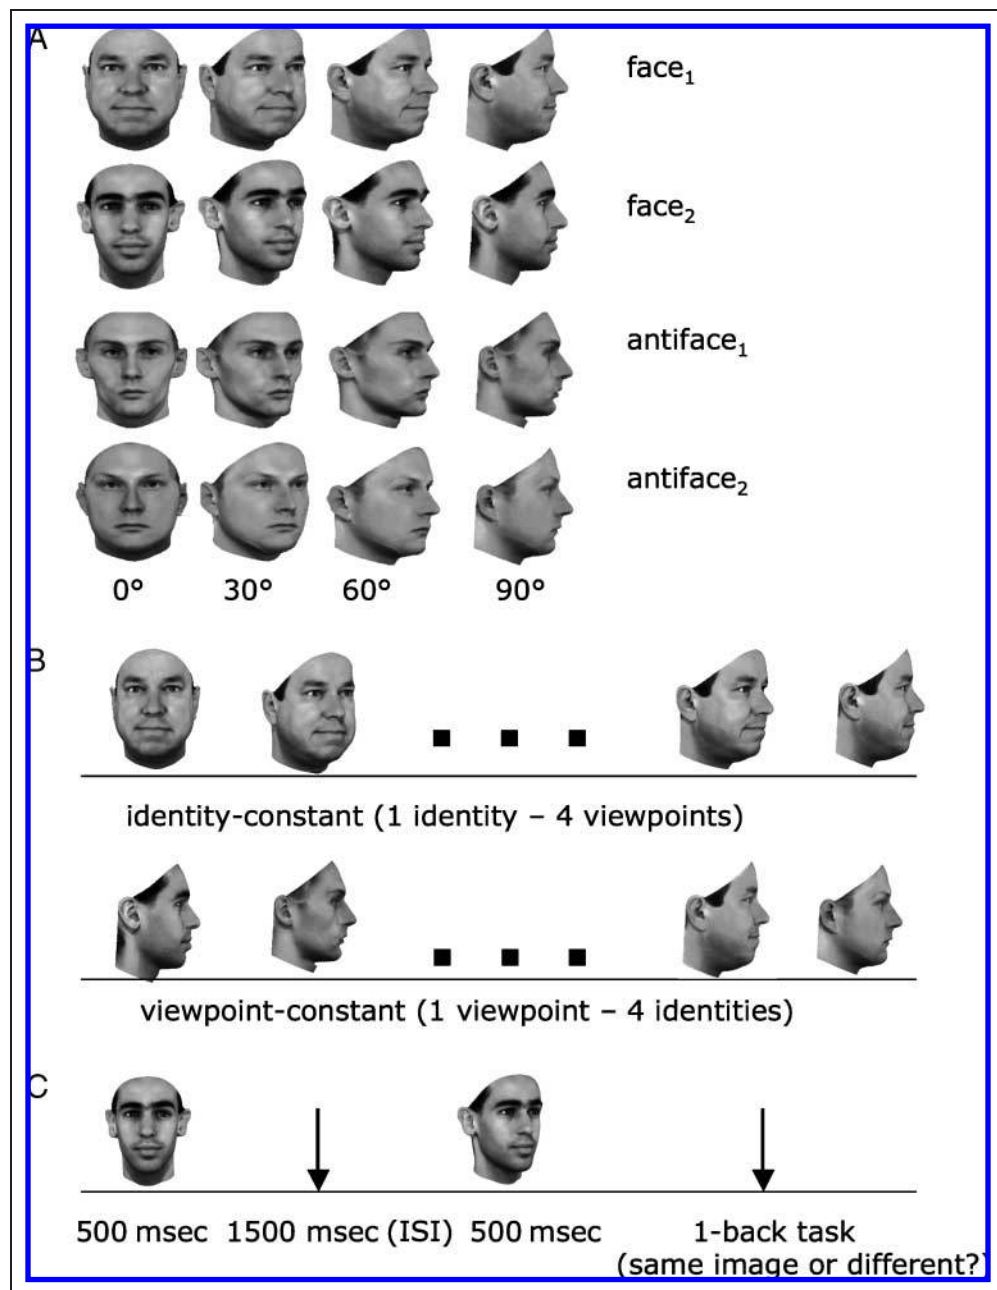
fMRI Experimental Protocol and Task

The experimental protocol consisted of a localizer session followed by an experimental session. The localizer session was used to find voxels in ventral temporal (VT) cortex that respond differentially to faces, objects, and scrambled images. The localizer stimuli consisted of gray-scale images of human faces, objects (chairs and bottles), and scrambled images (Yovel & Kanwisher, 2004, 2005; Haxby et al., 2001). None of these images appeared in the experimental sessions.

Localizer Session

In the localizer session, participants viewed the faces, objects, and scrambled images in a blocked procedure used

Figure 1. Experimental stimuli and protocol. (A) Stimulus set consisted of the four male faces (2 original faces and their antifaces) viewed from four viewpoints: frontal (0°) to profile (90°) in increments of 30° . (B) Each identity-constant block consisted of a single facial identity presented four times from the each viewpoint, with viewpoint randomized within the block. A viewpoint-constant block consisted of images of all four facial identities presented from a single viewpoint four times with the identity randomized within the block. (C) In each trial, a face image appeared for 500 msec, followed by a 1500-msec blank interstimulus interval. A second stimulus followed and a response is made in a 1-back task.



in previous studies (Yovel & Kanwisher, 2004, 2005; Haxby et al., 2001) with slight modifications. In each localizer session, participants viewed six replications of three consecutive 12-sec blocks. Each block contained 12 images of a single category, presented in a random order. Each image appeared for 200 msec, followed by an 800-msec blank interstimulus interval. The blocks were preceded and followed by 10 sec of fixation. Participants performed a 1-back task during the scan in which they were instructed to respond “same” or “different” image to the consecutively presented images. Participants responded “same” only when the exact same image followed the previous image. The 1-back task was performed to maintain attention inside the scanner.

Experimental Session

Prior to the experimental session, to visually familiarize participants with the stimuli, they viewed each of four facial identities from four viewpoints. Each face was labeled with a “name” and appeared on the computer screen for 5 sec. The identities were presented, in turn, with the four views varying from frontal to profile in order. This session was conducted outside of the scanner, just prior to the scan session.

The experimental data were collected during four replications of eight blocked conditions. The eight conditions consisted of four identity-constant blocks and four viewpoint-constant blocks (Figure 1B). In an *identity-constant* trial

block, 0°, 30°, 60°, and 90° views of a single facial identity were presented four times each in random order. In a *viewpoint-constant* trial block, images of the four facial identities from a single viewpoint were presented four times each in random order. A block lasted for 32 sec and was preceded and followed by a 10-sec fixation point. An image within a block appeared for 500 msec, followed by a 1500-msec blank interstimulus interval (Figure 1C). The image location on the screen was set randomly to one of eight locations to avoid effects of apparent motion. Participants performed a 1-back task in both the identity-constant and viewpoint-constant blocks. Participants responded “same” when the exact same image followed the previous image. To minimize confounds of repetition, the blocks and the individual identities and viewpoints within each block were presented in random order. The stimulus sequences were presented using E-Prime 1.1 (Psychological Software Tools, Pittsburgh, PA) using a Windows PC.

Subjects

Eight healthy subjects (4 men, age range = 20–45 years) with normal or corrected-to-normal vision volunteered to participate in the fMRI experiment. Participants gave written informed consent to participate in the experiment. The Institutional Review committees at the University of Texas at Dallas and the University of Texas Southwestern Medical Center at Dallas approved the experimental protocol.

Data Acquisition and Image Processing

Functional images were acquired on a 3-T MR system (Achieva; Philips Medical Systems, Best, The Netherlands) with an eight-channel SENSE head coil. A high-resolution (voxel size = $1 \times 1 \times 1$ mm) MP-RAGE structural scan was acquired prior to the functional scans. The blood oxygen level dependent signal was obtained with echoplanar imaging transverse images (TR = 2000 msec, TE = 30 msec, flip angle = 80°, FOV = 220 mm, 38 slices, voxel size = $3.44 \times 3.44 \times 4.00$ mm) that covered the entire cortex.

The localizer and experimental imaging data obtained for each participant were preprocessed using SPM5 (www.fil.ion.ucl.ac.uk/spm/software/spm5/). The volumes were corrected for slice timing, realignment, and coregistration using the default parameters in SPM5. The data from one of the eight subjects were eliminated from further analysis due to excessive head motion.

Voxel Selection for Input to Classifier

We loaded the preprocessed localizer and experimental datasets into Matlab. We selected only the voxels within the VT region whose activity varied significantly across the


three stimulus categories (faces, objects, and scrambled images), using the localizer data. An analysis of variance, with a p criterion $< .0001$ was performed on individual voxel activity to generate a functional mask of voxels for input to the classifier. This is a more liberal voxel selection process than what was used in previous studies. Figure 2 shows the neural regions localized in the VT voxel mask for one of the participants. For all participants, the VT mask included regions in and around fusiform gyrus and occipital face areas. The pSTS and the anterior temporal areas were localized for two of the participants. We will say more about the limited inclusion of pSTS shortly. In addition to these standard face-selective areas, the VT mask included lateral occipital areas for all participants. The voxels selected formed a *VT mask* with an average of 502.7 voxels ($SD = 110.3$) across the eight participants.

We noted that the VT masks did not consistently include pSTS, which may play a role in viewpoint discrimination vis-à-vis the connection between viewpoint and social attention (Pelphrey, Viola, & McCarthy, 2004; Haxby, Hoffman, & Gobbini, 2002; Haxby et al., 2000). The failure to locate pSTS is a common problem with localizers that use static face images to find face-selective cortex (Fox, Iaria, & Barton, 2009; Kanwisher et al., 1997). Thus, for some viewpoint classifications we report in this article, we selected voxel clusters in left and right pSTS based on the anatomical locations. To create the pSTS masks, we drew spherical ROIs around the right and left pSTS loci using anatomical landmarks from the human brain atlas as a guide (cf. Mai, Assheuer, & Paxinos, 1997). These were adjusted individually, as needed, to center the region at the posterior termination of STS. Each ROI contained 33 voxels (with a spherical radius of 8.5 mm). In all cases, we varied the radius of the spherical ROIs to verify stability of the classification.

Discrimination of Neural Response Patterns for Identity and Viewpoint

Pattern-based classifiers were implemented to measure the neural discriminability of all possible pairs of face identity (6 discriminations: face₁ vs. face₂, face₁ vs. antiface₁, face₁ vs. antiface₂, face₂ vs. antiface₁, face₂ vs. antiface₂, antiface₁ vs. antiface₂) and all possible pairs of viewpoint (0° vs. 30°, 0° vs. 60°, 0° vs. 90°, 30° vs. 60°, 30° vs. 90°, 60° vs. 90°). The pattern classifiers were implemented separately for each participant. We created two counterbalance conditions from different halves of the data using the odd and even runs of the experimental session (cf. O’Toole et al., 2005; Haxby et al., 2001). In each counterbalance condition, the training and test datasets contained 60 scans (30 per category) with n voxels per scan, where n equals the number of preselected voxels. The algorithm proceeded as follows. First, to create an abbreviated representation of the individual scans, we applied a principal component (PC) analysis to all scans from the training set. Individual scans were projected into the PC analysis space to

Figure 2. Multiple axial slices showing the neural regions included in the VT voxel mask for a single participant. The VT mask typically included fusiform gyrus, occipital face area, and lateral occipital cortex for all participants (the right hemisphere appears on the right for the brain images). All highlighted voxels were included in the VT mask.



determine their coordinates on each PC. The coordinates were then used to represent the individual scans for input to the classifier.

Because the individual PCs vary in their usefulness for discriminating the neural activation patterns that result from viewing the two categories of stimuli (e.g., face₁ vs. face₂), the next step was to select individual PCs for classification based on their usefulness using the training set. The prescreening process minimizes problems with overfitting to a particular training dataset. It is also useful for reducing noise in classifying the test set when dimensions unrelated to the experimental variable are included. We assessed the utility of the PCs by training a series of single-dimension linear discriminant classifiers using the coordinates of the scans on the individual PCs. In order to measure the performance of the individual PCs in these classifiers, we measured neural discriminability using the signal detection measure d' , measured as Z -score (hit rate) – Z -score (false alarm rate). For example, in discriminating the neural signals for face₁ versus face₂, the hit rate was defined as the proportion of face₁ patterns classified correctly and the false alarm rate was the proportion of face₂ patterns classified incorrectly. The use of d' corrects for classifier bias that may occur when there is overfitting (e.g., a classifier with a bias to categorize scans into a particular category). A threshold d' was set on the training data to select PCs to be combined into a higher dimensional classifier for classifying scans from the test data.

The result for each discrimination problem was a low-dimensional subspace classifier tailored to discriminating individual stimulus pairs by identity or viewpoint. For both the identity and viewpoint classifiers, we tested a range of thresholds to verify the stability of the pattern of discrimination and to find values that optimized accuracy. The d' threshold for inclusion in the viewpoint classifiers was 0.50. The d' threshold for inclusion in the identity classifiers was 0.75. This classification procedure was implemented for each counterbalance and for each participant. The classification results we report are based on averages over the participants and the two counterbalance runs.

Identity and Viewpoint Discrimination: Behavioral Experiments

As noted, we rely on stimulus-based perceptual predictions to constrain our interpretation of the neural classification data. Thus, we performed two behavioral experiments outside of the scanner to measure the perceptual discriminability of the face identities and viewpoints tested in the fMRI study. Naïve participants who had not taken part in the fMRI study were recruited from the subject pool at The University of Texas at Dallas. For the identity comparisons, on each trial, participants ($n = 6$) viewed a pair of images for 500 msec. Participants were asked to judge if the two images belonged to the “same” or “different”

person as quickly as possible. The viewpoint of the two faces was varied to include all possible view and identity pairings. Over the course of 192 trials, 96 trials consisted of all possible “different” pairs of identities across all possible viewpoints. The remaining 96 trials, consisted of 16 “same” identity pairs across all viewpoints repeated once, and the 16 pairs with the exact same images (i.e., same face from same viewpoint) repeated five times each. This was done to balance the number of “same” and “different” trials.

The trials were identical for the viewpoint discrimination experiment, but in this case, participants ($n = 8$) judged if the two images had the “same” or “different” viewpoint angle. Exposure time in this experiment was reduced to 200 msec based on pilot data indicating that the task was easier than the identity discrimination.

The behavioral experiments differed from the 1-back task done in the scanner because our goal was to measure perceptual rather than neural discrimination of the face identities (over viewpoint) and the face views (over identity change). The pattern classification analysis required blocked presentations of the identities or viewpoints. For the perceptual discrimination, participants made judgments about whether identity or viewpoint matched in simultaneously presented pairs of faces. Ultimately the goal was to determine how similar two face identities (or viewpoints) appeared. Thus, although the perceptual task differed from the neural task, it provided data that were analogous to the neural discrimination data.

RESULTS

Neural and Perceptual Discrimination of Identity

We obtained low to moderate d' scores for discriminating the neural activation patterns for most, but not all, participants. We eliminated two participants who had median d' values across the six neural discriminations at or less than zero. Figure 3A shows the average discrimination scores for the remaining five participants on the identity pairs. These are displayed in rank order of discriminability, that is, the most to least discriminable pair. Four of the six face pairs were discriminated at levels above chance. The face–antiface pairs were ranked second and third in discrimination order. The original unaltered faces ranked fourth. The best-discriminated pair was a face and antiface pair from different people.

Figure 3B shows the reaction times from the perceptual experiment for judging the identity pairs as “different.” These data indicate the level of difficulty for determining that two faces were different identities, again assessed over change in viewpoint. For ease of comparison, we plotted the reaction times for each identity pair, in order of the neural rankings, that is, the neural classifier’s best-discriminated pair to the worst-discriminated pair. Reaction times in the perceptual experiment for the pairs discriminated neurally at levels above chance were sig-

nificantly faster than reaction times for the pairs that were not neurally discriminated [$F(1, 25) = 15.03, p < .001$], indicating agreement between the neural and perceptual discriminability of the face pairs.

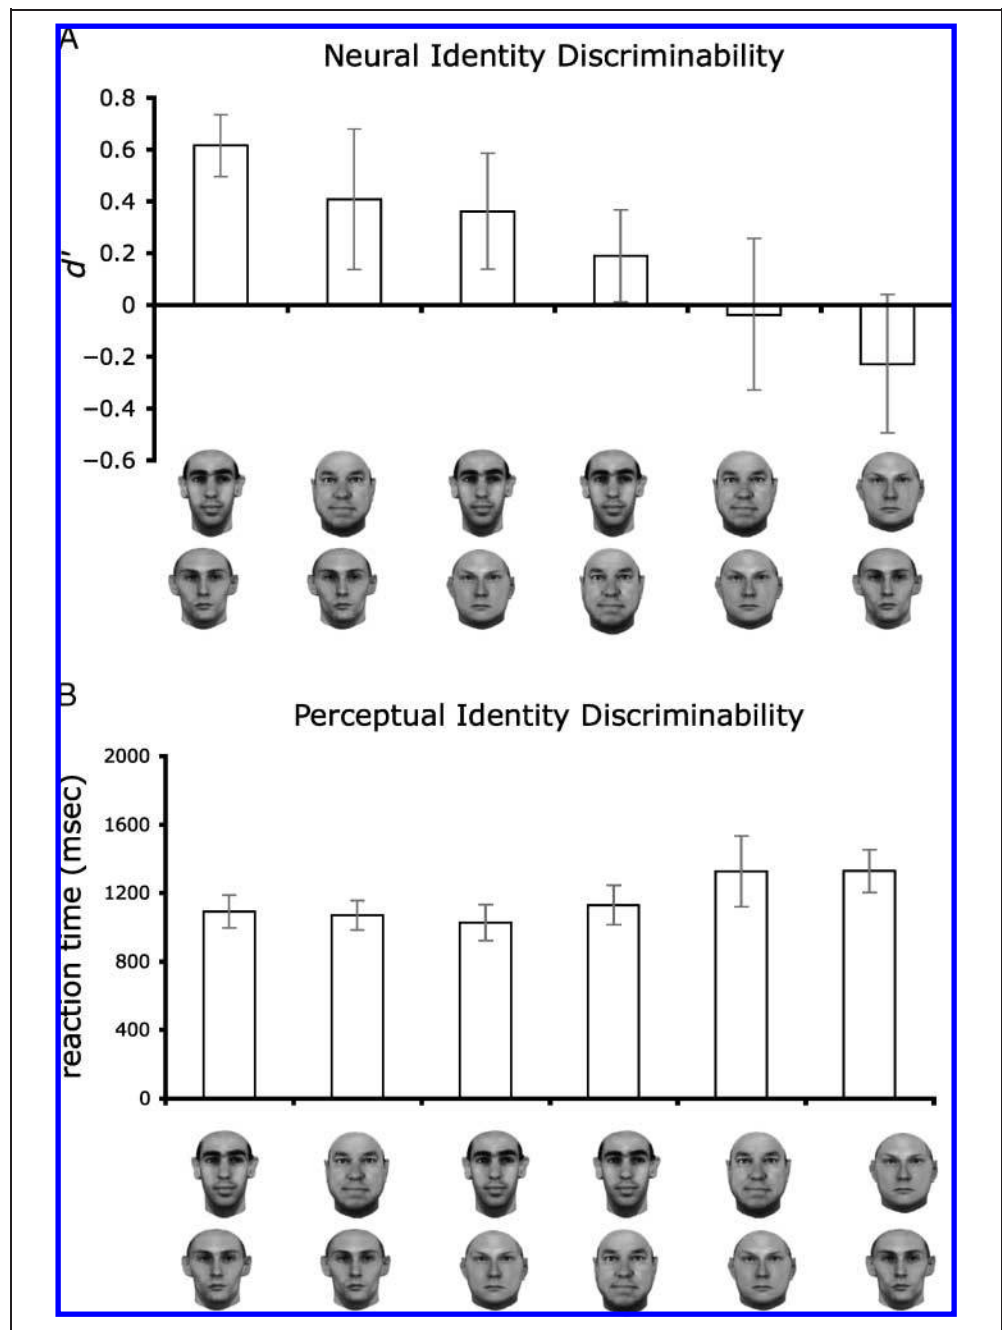
Based on indications from the fMR-A literature that face-selective areas such as FFA are minimally tolerant to viewpoint variation, we did not expect face-selective areas to support discrimination. For completeness, however, we tested neural discrimination of identity also with face-selective voxels. These voxels were found using the standard (face > object) contrast on the localizer session data. Consistent with the fMR-A findings, the pattern-based classifier failed to discriminate identity independent of viewpoint based on face-selective voxels alone.

Neural and Perceptual Discrimination of Viewpoint

We obtained low to moderate discrimination scores (d') for the neural activation patterns elicited in response to all possible pairs of viewpoints (0° vs. 30° , 0° vs. 60° , 0° vs. 90° , 30° vs. 60° , 30° vs. 90° , 60° vs. 90°) for six out of the seven participants (Figure 4A). One participant was eliminated for a median d' s across the six neural discriminations less than zero. Figure 4A shows above-chance neural discrimination for five of the six viewpoint pairs. We predicted that the viewpoint pairs with a small angular disparity would have lower discrimination scores as compared to pairs with a high angular disparity. We found no agreement between the pattern of neural discrimination and the pattern of perceptual discrimination. The results of the behavioral experiment appear in Figure 4B (all data) and Figure 4D (averaged over angular disparity conditions). These behavioral data support our original stimulus-driven prediction that larger viewpoint disparities give rise to more accurate viewpoint discriminations. More precisely, the perceptual data indicated that 30° angular disparities were discriminated less accurately than 60° changes [$F(1, 7) = 9.44, p < .05$] and that 60° angular disparities were discriminated less accurately than the 90° change [$F(1, 7) = 38.67, p < .01$] (Figure 4B and D). Thus, neural and perceptual discrimination do not agree.

The lack of the predicted pattern could be due to the fact that the individual VT voxel masks used by the classifiers did not include pSTS consistently for all subjects. This region may be involved in processing head orientation and gaze direction for social attention (Pelphrey et al., 2004; Haxby et al., 2000; Hoffman & Haxby, 2000). We therefore repeated the classification algorithm using a mask made by combining the VT mask with the left and right pSTS masks. Discrimination performance was similar to that found with the VT mask alone, again out of accord with the perceptual data. Next, we looked at the right and left pSTS masks without the VT mask, both in combination and separately. Only the right pSTS mask showed above-chance performance that was in partial accord with perceptual discrimination.

Figure 3. Neural and perceptual discrimination of identity. (A) Neural discrimination scores (d') for the six identity pairs averaged across five participants. The scores are displayed in rank order of discriminability from the best- to the worst-discriminated pair. Four of the six face pairs were discriminated at levels above chance. The face–antiface pairs were ranked second and third in discrimination order. The original unaltered faces ranked fourth. The best-discriminated pair was a face and antiface pair from different people. Error bars represent one standard error of the mean across participants. (B) Perceptual discriminability represented as reaction times for the six identity pairs averaged across six different participants. The reaction time for each identity pair is plotted in order of the neural rankings, that is, the neural classifier’s best- to worst-discriminated pair. Reaction times for the pairs discriminated neurally at levels above chance were significantly faster than reaction times for the pairs that were not neurally discriminated.



Viewpoint Discrimination Performance for Right pSTS

Figure 4C illustrates the discrimination scores for right pSTS averaged for the three angular changes (i.e., 30°, 60°, and 90°) across the six subjects from Figure 4A. The results of the classifier for right pSTS (Figure 4C) showed only one viewpoint condition above chance, 0° vs. 90°. This viewpoint difference is the largest one we tested and is also the viewpoint difference discriminated most accurately in the perceptual experiment (see Figure 4D). The standard error bars indicate that only the largest angular disparity condition was discriminated with the neural data at levels above chance. The previous re-

sults with the VT mask are plotted for comparison in Figure 4C. As can be seen, discrimination with the VT mask is more accurate, but is inconsistent with the pattern of perceptual discrimination.

DISCUSSION

The present study provides the first demonstration of the discriminability of neural response patterns for individual facial identity over substantial changes in viewpoint and for viewpoint over changes in identity. For identity, we found reliable neural discrimination that matched the perceptual discriminability using a broad span of VT cortex.

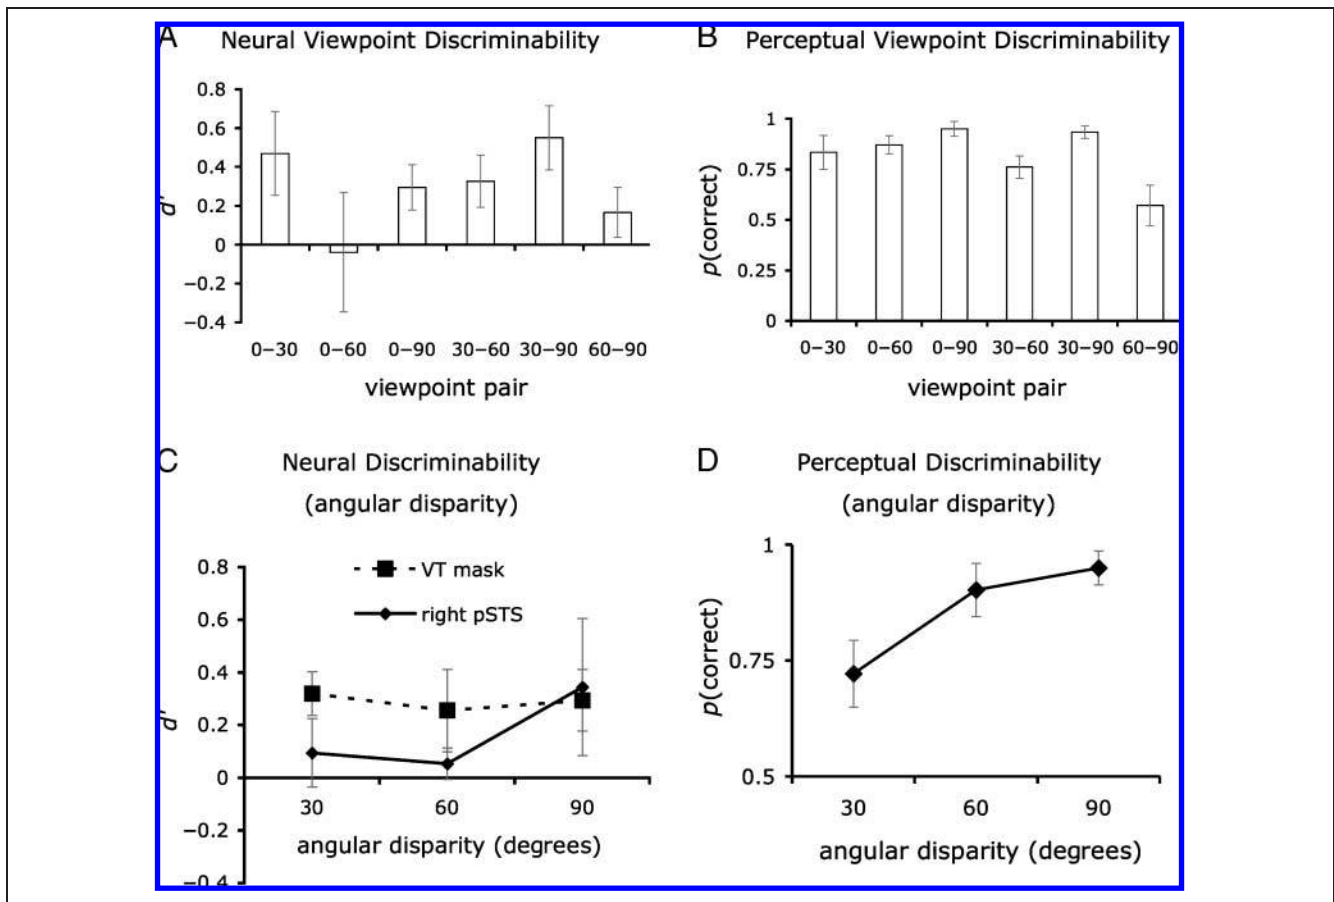


Figure 4. Neural and perceptual discrimination of viewpoint. (A) Neural discrimination scores (d') for the six viewpoint pairs averaged across six participants. Above-chance neural discrimination was obtained for five of the six pairs. (B) The perceptual and neural discrimination results disagree. (C) Neural discrimination for the VT mask from A, averaged over angular disparity, plotted with the discrimination data from right pSTS. The right pSTS results are more consistent with the perceptual predictions than the VT results. (D) For comparison, the perceptual discrimination for viewpoint from B, averaged over angular disparity.

The level of discrimination for individual identities was moderate to low, but was consistent across subjects and above chance for four of six face pairs. Relative to the performance reported previously for coarse-scale dissociations between faces and objects (O'Toole et al., 2005; Hanson et al., 2004; Carlson et al., 2003; Cox & Savoy, 2003; Spiridon & Kanwisher, 2002; Haxby et al., 2001), these moderate discrimination scores are in the range we expected. The fact that the code generalized across viewpoint indicates that the identity information tapped by the classifier transcends viewpoint-dependent image-based codes. The agreement between neural and perceptual discrimination is consistent with a code that is essentially high-level visual in nature. For viewpoint, we found reasonable levels of neural discrimination using a broad span of VT cortex. Additionally, we found some evidence for neural discrimination in accord with perceptual discrimination performance, but only in right pSTS.

The classifier in the present study generalized identity discrimination over viewpoint and ultimately required a broader area of cortex than the identity classification done by Kriegeskorte et al. (2007). In their study, the

classifier discriminated neural activation patterns elicited in response to viewing two face identities (a male and a female) pictured from the same viewpoint. Kriegeskorte et al. found individuation of identity in a single area of cortex (aIT). The discrimination task we did was at a more general level of face identity coding than the task done by Kriegeskorte et al. and may have required more complex visual processing.

The finding that the neural discrimination of individual identity over viewpoint change required a broad area of VT cortex is consistent with fMR-A and repetition suppression methods indicating minimal tolerance to viewpoint change in face-selective areas. Although it has been established that FFA codes features important for specifying face identity, the present study combined with the findings of adaptation-based studies indicates that traditionally defined face-selective areas operating independently cannot account for face recognition over viewpoint change. Indeed, two characteristics of a neural code useful for face recognition are high selectivity for identity-specific changes in face structure and an ability to generalize across image-based changes that are not relevant for

identification. Although high selectivity for identity-specific information has been demonstrated unequivocally in FFA, there is no evidence for generalization capacity beyond small viewpoint changes, even for familiar faces. The competing constraints of finely tuned sensitivity to identity and generality across viewing parameters define a complex computational problem that may involve a collaborative and interactive dialog between several high-level visual areas. The possibility that face recognition is done by cooperative and competitive interactions among brain regions is consistent with Bayesian framework for visual recognition that assumes active generative models of objects and faces in addition to feedforward visual processing mechanisms (Yuille & Kersten, 2006). A recent study of a prosopagnosic patient demonstrates an important role for LOC in processing faces (Dricot et al., 2008) and is likewise consistent with the collaborative computational framework in suggesting a complementary role for LOC to face-selective areas such as FFA and the occipital fusiform area. As noted, the more liberal voxel selection method we used included lateral occipital areas consistently across subjects, making it available as a resource for the identity classification.

We did not dissect the present results in terms of the individual contribution of functionally defined face- and object-selective brain areas for both theoretical and practical reasons. From the theoretical perspective, the questions we address here about face identity codes have been investigated thoroughly using adaptation-based methods that operate in predefined regions of interest. The results of these studies point consistently to the view-dependent nature of the face codes in face-selective areas. Additionally, previous studies have not considered the potential for interactions across multiple regions of interest, which might be part of a cooperative and competitive processing network for recognition.

From the practical perspective, based on our initial assumption that neural classification at this fine-grained scale of identity would be a challenging problem, our first goal was to achieve reliable classification. Given the literature showing that face-selective areas show limited viewpoint generalization and the concern that classification would require access to most or all of the relevant neural information, we began with a liberal criterion for voxel inclusion. For identity, we found discrimination and accord between neural and perceptual data at the level of individual stimulus pairs. Thus, by parsimony, further subdivision of the brain areas was unnecessary.

Although delineating a large area that discriminates between different faces may seem like a step backward in our understanding of the processes underlying the individuation of faces, it may be an important prerequisite step for making new progress on the problem. There is now sufficient evidence from adaptation studies to indicate that individual local areas, by themselves, are not likely to be capable of identity individuation that generalizes over changes in viewing conditions. It is perhaps

worth taking a step backward to reconsider the contribution of a larger area of cortex to this difficult task. If the ultimate solution to the problem involves a network of high-level visual areas, the next research steps should be aimed at determining which components contribute to the process and how the various components interact.

For viewpoint, the perceptual and neural data disagreed and so we examined pSTS and found evidence that 0° versus 90° viewpoints were discriminable in right pSTS. We are uncertain why we found moderate, but perceptually discordant, discrimination of viewpoint in broader VT cortex. One speculative possibility, consistent with Haxby et al. (2000), is that viewpoint information is coded both in ventral face areas and in pSTS, but with different functions. In the ventral stream, the goal of viewpoint processing would be to compensate or normalize these viewpoint variations to establish identity. In the pSTS stream, the goal would be an accurate estimate of head direction. These redundant codes are likely to be qualitatively different. Classification with the VT voxel mask may have been more accurate because it had access to some components of both codes, but perceptually discordant because the classification combined codes optimized for different functions. Thus, the finding that an area is sensitive to a stimulus dimension does not reveal the function for which the dimension is coded when there are multiple behavioral uses for the information. In these cases, the use of perceptual data to generate predictions for neural discrimination can constrain the interpretation of the neural findings. The technique of generating neural predictions from perceptual data has been used previously in single-cell studies (e.g., Op de Beeck, Wagemans, & Vogels, 2001; Young & Yamane, 1992) and in neuroimaging studies (Haushofer, Livingstone, & Kanwisher, 2008; O'Toole et al., 2005; Edelman, Grill-Spector, Kushnir, & Malach, 1998). In both cases, a supportive link between neural response patterns and perception can anchor the interpretation of neural data, especially in cases where redundant neural codes with different functions may exist.

Finally, it should be noted that the fact that it was possible to discriminate the neural responses elicited by individual identities over viewpoint change is *only* a prerequisite step for understanding face identity coding in cortex. Our finding that discrimination required a relatively broad area of VT cortex does not preclude the possibility that face identity can be decoded in face-selective areas. Evidence for this, however, may require far better spatial resolution than is available currently with fMRI, and the question may ultimately be decided only with methods that can operate over a broad area of cortex at the resolution of individual neurons. For present purposes, our findings support the following conclusions. For identity, we found information available across VT cortex to support face identity discrimination at a level that transcends viewpoint-dependent image-based codes. Moreover, the agreement between neural and perceptual

discrimination indicates that the classification was based on a high-level visual code. For viewpoint, we conclude that there is information available across VT cortex to decode viewpoint. The lack of accord between the perceptual and neural discrimination for the broader VT cortex, in combination with the limited accord for right pSTS, suggests the possibility of redundant viewpoint codes that may be qualitatively and spatially distinct. Overall, these results highlight the importance of applying perceptual constraints to the interpretation of functional neuroimaging data.

Acknowledgments

We thank the Advanced Imaging Research Center at the University of Texas Southwestern Medical School, Dallas, TX for support of scan time. We also thank Christina Wolfe for data analysis support and figure preparation. Thanks are also due to two anonymous reviewers who provided helpful comments on a previous version of the manuscript.

Reprint requests should be sent to Vaidehi S. Natu, School of Behavioral and Brain Sciences, GR4.1, The University of Texas at Dallas, Richardson, TX 75080, or via e-mail: vsnatu@utdallas.edu.

Notes

1. Fang et al. (2006) distinguish between view-tuned and viewpoint-sensitive in a way that does not map directly onto the general definition of viewpoint-dependent used in most studies.
2. One of the eight subjects was personally familiar with the original faces.

REFERENCES

- Andrews, T. J., & Ewbank, M. P. (2004). Distinct representations for facial identity and changeable aspects of faces in the human temporal lobe. *Neuroimage*, *23*, 905–913.
- Barton, J. J. S., Press, D. Z., Keenan, J. P., & O'Connor, M. (2002). Lesions of the fusiform face area impair perception of facial configuration in prosopagnosia. *Neurology*, *58*, 71–78.
- Blanz, V., & Vetter, T. (1999). A morphable model for the synthesis of 3D faces. In *SIGGRAPH'99: Proceedings of the 26th annual conference on computer graphics and interactive techniques* (pp. 187–194). New York: ACM Press/Addison-Wesley.
- Carlson, T. A., Schrater, P., & He, S. (2003). Patterns of activity in the categorical representations of objects. *Journal of Cognitive Neuroscience*, *15*, 704–717.
- Cox, D., & Savoy, R. (2003). Functional magnetic resonance imaging (fMRI) “brain reading”: Detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage*, *19*, 261–270.
- Damasio, A. R., Damasio, H., & Van Hoesen, G. W. (1982). Prosopagnosia: Anatomical basis and neurobehavioral mechanism. *Neurology*, *32*, 331–341.
- Dricot, L., Sorger, B., Schiltz, C., Goebel, R., & Rossion, B. (2008). The roles of “face” and “non-face” areas during individual face discrimination: Evidence by fMRI adaptation in a brain-damaged prosopagnosic patient. *Neuroimage*, *40*, 318–332.
- Edelman, S., Grill-Spector, K., Kushnir, T., & Malach, R. (1998). Toward direct visualization of the internal shape representation space by fMRI. *Psychobiology*, *26*, 309–321.
- Eger, E., Ashburner, J., Haynes, J. D., Dolan, R. J., & Rees, G. (2008). fMRI activity patterns in human loc carry information about object exemplars within category. *Journal of Cognitive Neuroscience*, *20*, 356–370.
- Eger, E., Schweinberger, S. R., Dolan, R. J., & Henson, R. N. (2005). Familiarity enhances invariance of face representations in human ventral visual cortex: fMRI evidence. *Neuroimage*, *26*, 1128–1139.
- Ewbank, M. P., & Andrews, T. J. (2008). Differential sensitivity for viewpoint between familiar and unfamiliar faces in human visual cortex. *Neuroimage*, *40*, 1857–1870.
- Fang, F., Murray, S. O., & He, S. (2006). Duration-dependent fMRI adaptation and distributed viewer-centered face representation in human visual cortex. *Cerebral Cortex*, *17*, 1402–1411.
- Fox, C. J., Iaria, G., & Barton, J. (2009). Defining the face-processing network: Optimization of the functional localizer in fMRI. *Human Brain Mapping*, *30*, 1637–1651.
- Fox, C. J., Moon, S.-Y., Iaria, G., & Barton, J. S. (2009). The correlates of subjective perception of identity and expression in the face network: An fMRI adaptation study. *Neuroimage*, *44*, 569–580.
- Ganel, T., & Goshen-Gottstein, Y. (2004). Effects of familiarity on the perceptual integrality of the identity and expression of faces: The parallel-route hypothesis revisited. *Journal of Experimental Psychology: Human Perception and Performance*, *30*, 583–597.
- Gilaie-Dotan, S., & Malach, R. (2007). Sub-exemplar shape tuning in human face-related areas. *Cerebral Cortex*, *17*, 325–338.
- Grill-Spector, K., Knouf, N., & Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nature Neuroscience*, *7*, 555–562.
- Grill-Spector, K., Kushnir, T., Hendler, T., Edelman, S., Itzhack, Y., & Malach, R. (1999). Differential processing of objects under various viewing conditions in human lateral occipital complex. *Neuron*, *24*, 187–203.
- Hadjikhani, N., & De Gelder, B. (2002). Neural basis of prosopagnosia: An fMRI study. *Human Brain Mapping*, *16*, 176–182.
- Hancock, P. J. B., Bruce, V., & Burton, A. M. (2000). Recognition of unfamiliar faces. *Trends in Cognitive Sciences*, *4*, 330–337.
- Hanson, S. J., Matsuka, T., & Haxby, J. V. (2004). Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: Is there a face area? *Neuroimage*, *23*, 156–166.
- Haushofer, J., Livingstone, M., & Kanwisher, N. (2008). Multivariate patterns in object-selective cortex dissociate perceptual and physical similarity. *PLoS Biology*, *29*, 1459–1467.
- Haxby, J., Hoffman, E., & Gobbini, M. (2002). Human neural systems for face recognition and social communication. *Biological Psychiatry*, *51*, 59–67.
- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representation of faces and objects in ventral temporal cortex. *Science*, *293*, 2425–2430.
- Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences*, *4*, 223–233.
- Henson, R. N., Shallice, T., & Dolan, R. J. (2000). Neuroimaging evidence for dissociable forms of repetition priming. *Science*, *287*, 1269–1272.
- Hoffman, E. A., & Haxby, J. V. (2000). Distinct representations of eye gaze and identity in the distributed human neural system for face perception. *Nature Neuroscience*, *3*, 80–84.

- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, *17*, 4302–4311.
- Kriegeskorte, N., Formisano, E., Sorger, B., & Goebel, R. (2007). Individual faces elicit distinct response patterns in human anterior temporal cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *104*, 20600–20605.
- Leopold, D. A., O’Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, *4*, 89–94.
- Mai, J. K., Assheuer, J., & Paxinos, G. (1997). *Atlas of the human brain*. San Diego: Academic Press.
- Op de Beeck, H., Wagemans, J., & Vogels, R. (2001). Inferotemporal neurons represent low-dimensional configurations of parameterized shapes. *Nature Neuroscience*, *4*, 1244–1252.
- O’Toole, A. J., Jiang, F., Abdi, H., & Haxby, J. V. (2005). Partially distributed representations of objects and faces in ventral temporal cortex. *Journal of Cognitive Neuroscience*, *17*, 580–590.
- Pageler, N. M., Menon, V., Merin, N. M., Eliez, S., Brown, W. E., & Reiss, A. L. (2003). Effect of head orientation on gaze processing in fusiform gyrus and superior temporal sulcus. *Neuroimage*, *20*, 318–329.
- Pelphrey, K. A., Viola, R. J., & McCarthy, G. (2004). When strangers pass: Processing of mutual and averted social gaze in the superior temporal sulcus. *Psychological Science*, *15*, 598–603.
- Pourtois, G., Schwartz, S., Seghier, M. L., Lazeyras, F., & Vuilleumier, P. (2005a). Portraits or people? Distinct representations of face identity in the human visual cortex. *Journal of Cognitive Neuroscience*, *17*, 1043–1057.
- Pourtois, G., Schwartz, S., Seghier, M. L., Lazeyras, F., & Vuilleumier, P. (2005b). View-independent coding of face identity in frontal and temporal cortices is modulated by familiarity: An event-related fMRI study. *Neuroimage*, *24*, 1214–1224.
- Pourtois, G., Schwartz, S., Spiridon, M., Martuzzi, R., & Vuilleumier, P. (2009). Object representations for multiple visual categories overlap in lateral occipital and medial fusiform cortex. *Cerebral Cortex*, *19*, 1806–1819.
- Rotshtein, P., Henson, R. N. A., Treves, A., Driver, J., & Dolan, R. J. (2005). Morphing Marilyn into Maggie dissociates physical and identity face-representations in the brain. *Nature Neuroscience*, *8*, 107–113.
- Spiridon, M., & Kanwisher, N. (2002). How distributed is visual category information in human occipito-temporal cortex? An fMRI study. *Neuron*, *35*, 1157–1165.
- Troje, N. F., & Bühlhoff, H. H. (1996). Face recognition under varying pose: The role of texture and shape. *Vision Research*, *36*, 1761–1771.
- Valentin, D., Abdi, H., & O’Toole, A. J. (1994). Categorization and identification of human face images by neural networks: A review of linear auto-associator and principal component approaches. *Journal of Biological Systems*, *2*, 413–429.
- Vetter, T., & Troje, N. F. (1997). Separation of texture and shape in images of faces for image coding and synthesis. *Journal of the Optical Society of America*, *14*, 2152–2161.
- Young, M. P., & Yamane, S. (1992). Sparse population coding of faces in the inferotemporal cortex. *Science*, *256*, 1327–1331.
- Yovel, G., & Kanwisher, N. (2004). Face perception; domain specific, not process specific. *Neuron*, *44*, 889–898.
- Yovel, G., & Kanwisher, N. (2005). The neural basis of the behavioral face-inversion effect. *Current Biology*, *15*, 2256–2262.
- Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: Analysis by synthesis? *Trends in Cognitive Sciences*, *10*, 301–308.