Research Article

Probing the Visual Representation of Faces With Adaptation

A View From the Other Side of the Mean

Fang Jiang,[1] Volker Blanz,[2] and Alice J. O'Toole[1]

[1]School of Behavioral and Brain Sciences, The University of Texas at Dallas, and [2]Max-Planck-Institut für Informatik, Saarbrücken, Germany

ABSTRACT—*Sensory adaptation and visual aftereffects have long given insight into the neural codes underlying basic dimensions of visual perception. Recently discovered perceptual adaptation effects for complex shapes like faces can offer similar insight into high-level visual representations. In the experiments reported here, we demonstrated first that face adaptation transfers across a substantial change in viewpoint and that this transfer occurs via processes unlikely to be specific to faces. Next, we probed the visual codes underlying face recognition using face morphs that varied selectively in reflectance or shape. Adaptation to these morphs affected the perception of "opposite" faces both from the same viewpoint and from a different viewpoint. These results are consistent with high-level face representations that pool local shape and reflectance patterns into configurations that specify facial appearance over a range of three-dimensional viewpoints. These findings have implications for computational models of face recognition and for competing neural theories of face and object recognition.*

How does the brain encode the complex structure of human faces? A complete answer to this question includes both the feature dimensions underlying the neural encoding of faces and the visual information preserved in these codes. Adaptation has been used traditionally as a tool for probing the way basic visual dimensions such as color, motion, and orientation are encoded neurally. Recent studies show that it is possible also to probe the representations of higher-level visual forms, such as faces and complex shapes, using adaptation (Leopold, O'Toole, Vetter, & Blanz, 2001; Moradi, Koch, & Shimojo, 2005; Rhodes, Jeffery, Watson, Clifford, & Nakayama, 2003; Suzuki & Cavanagh, 1998; Watson & Clifford, 2003; Webster & MacLin, 1999; Webster, Kaping, Mizokami, & Duhamel, 2004; Zhao & Chubb, 2001). For example, people identify a face more accurately if they see it immediately after viewing its synthetically created *antiface*—a face with opposite features (Leopold et al., 2001; Moradi et al., 2005; see Fig. 1). This antiface adaptation is powerful enough to elicit an afterimage of the original face during a brief subsequent presentation of the average face (Leopold et al., 2001; Moradi et al., 2005).

Opponent-identity adaptation originates within perceptually salient feature dimensions of faces (Webster & MacLin, 1999; Webster et al., 2004). For example, selective adaptation to the configural elements of face shape (Webster & MacLin, 1999) and to the natural categorical dimensions of gender, ethnicity, and facial expression (Webster et al., 2004) facilitates the perception of faces with opposite values on these features. Opponent adaptation effects suggest that face encodings are based on contrastive neural mechanisms defined relative to average or neutral values of the feature dimensions that represent faces (Leopold et al., 2001).

In the present study, we extended the use of adaptation to investigate the nature of the underlying visual representation of facial identity at levels of neural processing that can support view-independent recognition. Previous studies have shown that face adaptation survives two-dimensional affine transformations, such as changes in retinal size and location (Leopold et al., 2001; Rhodes et al., 2003; Watson & Clifford, 2003; Zhao &

Address correspondence to Alice O'Toole, School of Behavioral and Brain Sciences, GR4.1, The University of Texas at Dallas, Richardson, TX 75083-0688, e-mail: otoole@utdallas.edu.

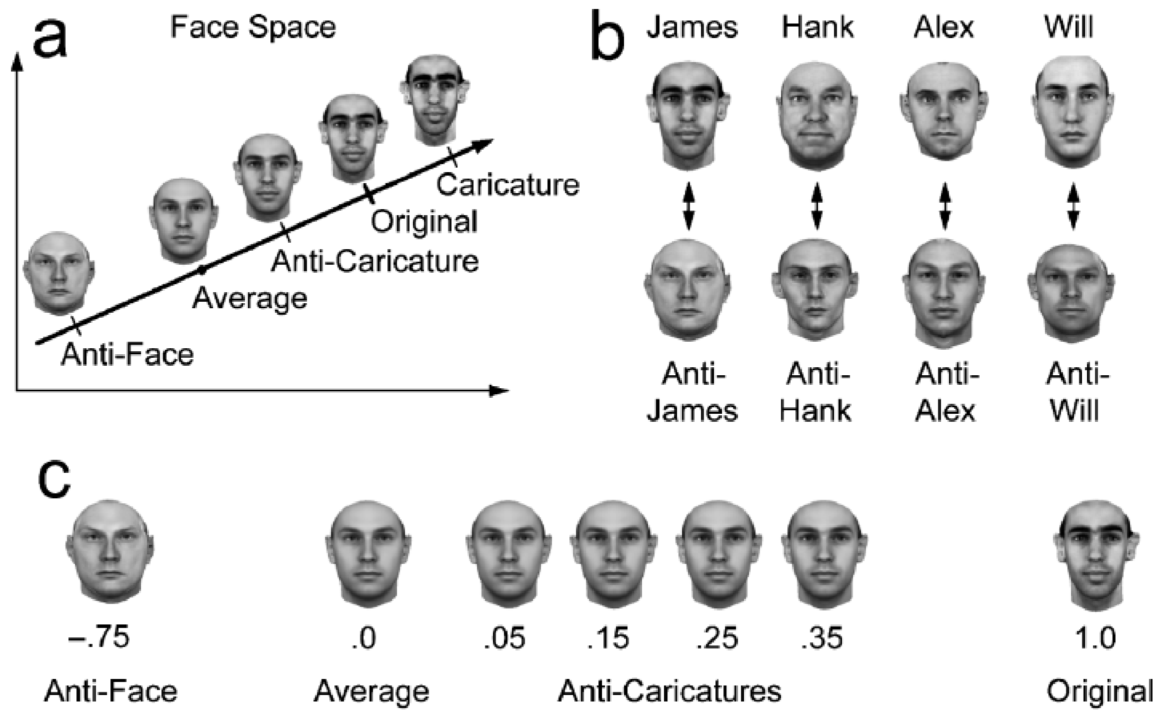

Fig. 1. Illustration of face space (a) and the stimuli used in Experiment 1 (b, c). In the multidimensional vector-space of three-dimensional faces, each point encodes the three-dimensional shape and surface reflectance (albedo) of a face. Morphing along the line from the average to an original face covers the range of anticaricatures, as it increases the identity strength. Further increasing identity strength beyond the original creates caricatures. Observers tend to perceive the faces as the same individual, which motivates the concept of an *identity trajectory*. The antiface of an original scan is located on the other side of the average. The illustrations in (b) show the four original scans that were used in the experiment and their antifaces. The faces in (c) illustrate identity strengths from the experiment, scaled in units of the distance of the original (1.0) from the average (0.0). Using antifaces with identity strengths equal to −0.75 avoided some morphing artifacts that tend to occur at more extreme values.

Chubb, 2001). The locus of adaptation effects is therefore likely to be in higher-level visual areas beyond those with strict retinotopic organization (Leopold et al., 2001). However, higher-level visual areas such as inferotemporal cortex or human lateral occipital cortex respond robustly across transformations that involve the three-dimensional structure of objects, such as changes in viewpoint (Booth & Rolls, 1998; Kourtzi, Erb, Grodd, & Bülthoff, 2003; Logothetis, Pauls, & Poggio, 1995; Vuilleumier, Henson, Driver, & Dolan, 2002) and illumination (Hietanen, Perrett, Oram, Benson, & Dittrich, 1992; Vogels & Biederman, 2002).

The intermixing of neurons with view-specific and view-independent receptive fields in high-level visual areas (Booth & Rolls, 1998) is consistent with hierarchical models of object recognition. These models achieve view-independent recognition using a hierarchy of neural mechanisms with view-dependent responses (Riesenhuber & Poggio, 1999, 2002). They accomplish this viewpoint-generalization task by learning to associate representations of the two-dimensional layout of features in different views of the same object.

The localization of face adaptation effects in higher-level visual areas is not proof that these effects are specific to faces, or that they involve mechanisms responsible for human face expertise. Indeed, robust adaptation effects have been reported for upright and inverted faces when the adapting and test face are in similar orientations—though the relative magnitude of adaptation for upright versus inverted faces varies across studies (Leopold et al., 2001; Rhodes et al., 2004; Webster & MacLin, 1999). Webster and MacLin (1999) concluded that although face adaptation strongly affects the perception of faces, it may not reflect processes specific to face recognition. Rather, they suggested that faces may be particularly useful for studying the effects of adaptation on form perception, because observers are highly sensitive to faces' configural properties.

The demonstrated adaptation to both upright and inverted faces is consistent with an adaptation locus in the human fusiform face area (FFA; Kanwisher, McDermott, & Chun, 1997). Paradoxically, the FFA responds with nearly equal strength to upright and inverted faces (Aguirre, Singh, & D'Esposito, 1999; Haxby et al., 1999; Kanwisher, Tong, & Nakayama, 1998; cf. Haxby, Hoffman, & Gobbini, 2000, for a discussion of these results)—despite evidence that it plays a role in human expertise for faces (e.g., Golby, Gabrieli, Chiao, & Eberhardt, 2001). The effectiveness of different kinds of stimuli for adapting facial-identity perception bears on the ongoing debate concerning whether the underlying representation of faces, presumably centered in the FFA, is modular (Kanwisher et al., 1997) or is based on a distributed network of shared feature codes for face and object perception (Haxby et al., 2001). The modular account suggests that altering the perception of facial identity via adaptation should require a face as the adapting stimulus. The distributed account suggests that the human representation of faces consists of neural mechanisms that respond to a conglomeration of shape- and reflectance-based features and their configurations. Therefore, stimuli that preserve the underlying visual components of these entities, but are not valid faces, should be effective in adapting the perception of facial identity.

In the present study, we used adaptation to probe the underlying nature of visual representations for face recognition within and across viewpoint. In the first experiment, we replicated the face-identity adaptation effect within the frontal viewpoint and extended it to show that adaptation transfers across a substantial change in three-dimensional viewpoint. We also examined the extent to which viewpoint transfer could be accounted for by adaptation to the general visual components of the face stimulus. The perception of facial identity was assessed, therefore, following adaptation to a nonface stimulus consisting of features from the frontal face arranged in a configuration consistent with the rotated face. This stimulus retained much of the diagnostic information specifying facial identity, but was inconsistent with a valid image of any three-dimensional face.

GENERAL METHOD

Stimuli

For all the experiments, the development of stimuli was based on the concept of a face space, which is common to psychological (Valentine, 1991), computational (Phillips, Moon, Rizvi, & Rauss, 2000), and neural (Young & Yamani, 1992) models of face recognition. For the present study, stimulus images were generated using a three-dimensional morphing program (Blanz & Vetter, 1999) that implemented a multidimensional face-space based on laser scans from 200 human heads.

Each point in this face space represents a natural human face. *Caricatures*, or exaggerated versions of a face, are made by morphing the face away from the average face along a straight line in the multidimensional space (Fig. 1). *Anticaricatures*, or less distinct versions of a face, are made by morphing the face toward the average face. The process of facial caricaturing and anticaricaturing results in a change in the perceived distinctiveness, but not the identity, of a face. We assume, therefore, that an individual is represented by an *identity trajectory* that originates at the average face (identity strength = 0) and is directed toward the position of the face in the space (identity strength = 1.0). *Antifaces*, or opposites of a face, lie on the other side of the mean in the face space (identity strength < 0; Blanz, O'Toole, Vetter, & Wild, 2000). An antiface is created by morphing a face backward along its identity trajectory, continuing beyond the average to the inverted side of the identity trajectory (Blanz et al., 2000; Leopold et al., 2001). Moving a face to the other side of the mean inverts its feature values on all axes of the multidimensional space. An antiface appears radically different from the original face, but still looks like a natural face. To ensure that the antifaces we used retained the sex of the original faces, we used the same-sex average face as the reference point in the space when we created the antiface for each original face.

Because the stimuli were generated with a three-dimensional model, it was possible to create images of view-changed faces. In this experiment, *view-changed faces* were rotated 30° to the right. Additionally, for Experiment 1, we created a set of antifaces with the internal features warped to the right side of the face, so that the features occupied positions (in the frontal view of the face) that were coincident with their positions in the view-changed version of the original face (Fig. 2b). The warping was accomplished with the liquefy tool in Photoshop CS®. These *warped faces* were grotesque and unnatural, but preserved much of the diagnostic information about facial identity. Specifically, the local shape and reflectance of the internal features were consistent with a frontal view of the face, and the configuration of features was consistent with a 30° rotated view.

Participants

A total of 128 undergraduates participated in the three experiments. The data from 5 participants were discarded because of poor performance across all caricature levels, leaving 53 participants in Experiment 1, 40 in Experiment 2, and 30 in Experiment 3. The number of participants in each adaptation condition varied between 10 and 15.

EXPERIMENT 1

Procedure

The experiment began with a short training session. During training, participants learned names for four original faces and then learned to associate these names to the 0.35-level anticaricatures of the faces. Feedback was given during training trials. During an additional 32 practice trials, participants identified the 0.35-level of anticaricatures, with or without adaptation, so that they would be further familiarized with the test protocol. These trials were excluded from the analysis.

For the experimental trials, frontal views of anticaricatures or the average face were presented briefly (200 ms), and participants' task was to identify each as one of the four original faces (Fig. 1). The anticaricatures varied in identity strength in equal steps (the specific levels used in this and the other experiments are indicated in the figures presenting results). Participants in the no-adaptation condition performed the face-identification task without adaptation. We refer to the trials in this condition as *anticaricature* and *neutral* trials, according to whether the test stimulus was an anticaricature or the average face, respectively. In the adaptation conditions, participants identified faces following 5 s of adaptation to frontal views of antifaces (within-viewpoint adaptation), view-changed antifaces (across-

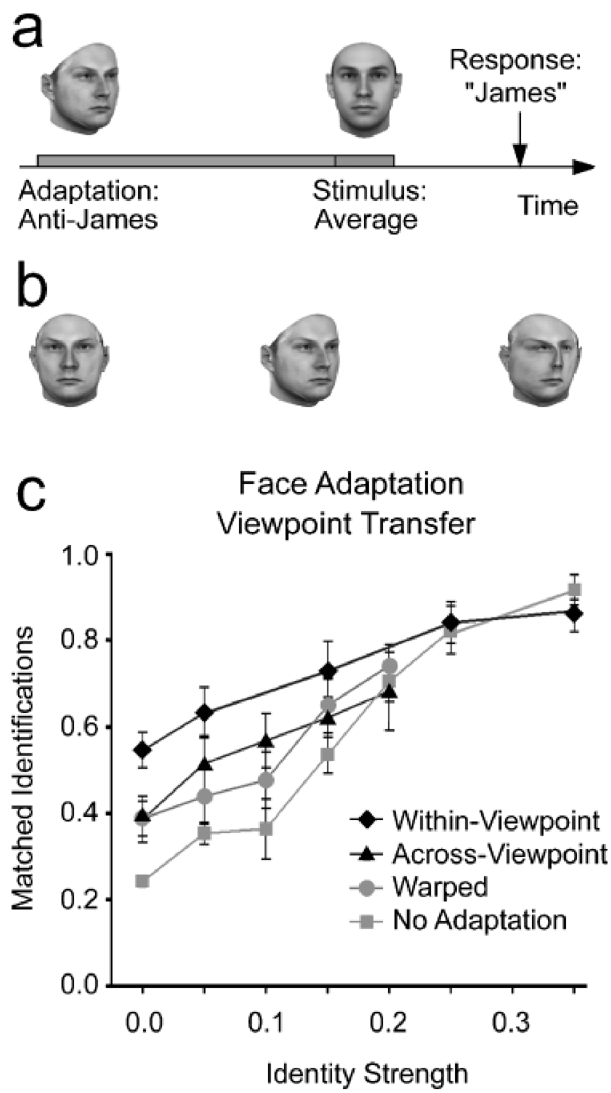

Fig. 2. Example stimuli and results from Experiment 1. The example trial (a) is taken from the across-viewpoint, neutral condition with adaptation. An antiface adapting stimulus appeared for 5 s, followed by the average face for 200 ms. Participants named the second stimulus, which was undefined in this case. The examples of stimuli in (b) include an antiface adapting stimulus as presented in the within-viewpoint condition (left), the corresponding antiface adapting stimulus in the view-changed condition (middle), and the corresponding warped antiface (right). The graph (c) shows the proportion of trials in the within-viewpoint, across-viewpoint, and warped conditions on which the test face was identified as the match to the adapting antiface stimulus, as a function of the identity strength of the test face. The no-adaptation line indicates the proportion of correct identifications for the faces as a function of identity strength (excluding the 0-identity-strength point, which reflects guessing responses to the average).

viewpoint adaptation), or warped antifaces (warped adaptation). The adaptation conditions included three types of trials: match, nonmatch, and neutral. On *match* trials, the test face was an anticaricature that was collinear along the same identity trajectory as the antiface adapting stimulus, so that the two faces straddled the mean. On *nonmatch* trials, the test face was an anticaricature from a different identity trajectory than the adapting antiface; we included these trials to prevent subjects from learning pairs of antifaces and anticaricatures. On *neutral* trials, the average face was presented for identification (see Fig. 2a for an example). For neutral trials with no adaptation, identification judgments would be expected to be distributed equally among the four faces. With adaptation, however, identification of the average face would be expected to be biased toward the face that matched (i.e., was on the same identity trajectory as) the antiface adaptation stimulus.

For the adaptation conditions, equal numbers of match, nonmatch, and neutral trials were presented, resulting in a total of 144 trials. Equal numbers of neutral trials and anticaricature trials were presented in the no-adaptation condition, resulting in 96 trials. Trial order was randomized individually for each participant.

Results

Adaptation effects were found both within and across viewpoint (Fig. 2c). The results for within-viewpoint adaptation replicate previous findings (Leopold et al., 2001)—adaptation improved anticaricature identification when the adapting and test stimulus were presented from the same viewpoint, $F(1, 18) = 23.77$, $p_{rep} = .996$, $\eta^2 = .328$, for identity strengths less than 0.25. The average face was identified as the match to the antiface adapting stimulus on .55 of the trials, significantly greater than the .25 expected by chance, $t(9) = 7.52$, $p_{rep} = .996$, $d = 2.38$. Adaptation to view-changed antifaces also facilitated the perception of matched, frontally viewed anticaricatures, $F(1, 20) = 6.36$, $p_{rep} = .93$, $\eta^2 = .13$, for identity strengths less than 0.20, albeit to a lesser extent than adaptation to frontal views of antifaces. The average face was identified as the match to the adapting stimulus on .39 of the trials, again significantly more often than expected by chance, $t(10) = 2.70$, $p_{rep} = .93$, $d = 0.82$.

Adaptation to warped antifaces likewise facilitated the perception of matched, frontally viewed anticaricatures, $F(1, 20) = 5.53$, $p_{rep} = .908$, $\eta^2 = .09$, for identity strengths less than 0.20. The magnitude of this adaptation effect did not differ significantly from that found for view-changed antifaces, $F(1, 20) = 0.21$, $p_{rep} = .393$, $\eta^2 = .006$, for identity strengths less than 0.20. Indeed, the average face was identified as the match to the antiface adapting stimulus on .39 of the trials, which is comparable to the proportion in the view-changed adaptation condition; again, this proportion was significantly greater than the .25 expected by chance, $t(10) = 3.60$, $p_{rep} = .97$, $d = 1.09$.

Discussion

The results indicate that adaptation to the identity of a face transfers across a substantial change in viewpoint, suggesting a locus of adaptation that reaches representations that can support view-transferable face recognition. The finding that warped faces produced identity adaptation of a magnitude similar to that produced by view-changed faces, however, suggests that the processes by which adaptation transfers across view change are not entirely specific to faces. Indeed, facial-identity perception was altered with comparable effectiveness using an adaptation stimulus that was not a real face, but preserved many of the perceptual features useful for face identification. This finding is consistent with a distributed view of high-level representations of faces, by which face-identity codes are accessible to the perceptual features from which they are composed.

In the next experiment, we considered explicitly the role of face shape and reflectance in identity adaptation. To do so, we measured the adaptation effect with faces that contained identity-specific information exclusively in either their shape or their reflectance.

EXPERIMENT 2

Method

Stimuli

A laser scan of a human head includes data on both three-dimensional shape (x, y, z) and surface reflectance (r, g, b). In the morphing program we used for these experiments (Blanz & Vetter, 1999), separate face vectors for three-dimensional shape and surface reflectance are constructed from 75,000 surface points sampled uniformly using dense point-to-point correspondence with a reference scan. Using this algorithm, for Experiment 2 we morphed three-dimensional shape and surface-reflectance data from individual faces separately to create faces that deviated from the average face only in three-dimensional shape or in surface reflectance (Fig. 3). A *shape-varying face* had the reflectance of the average face mapped onto the original face shape. A *reflectance-varying face* had its own reflectance mapped onto the shape of the average face.

Shape- and reflectance-varying anticaricatures and antifaces were created by morphing only the shape or reflectance data, respectively, along a face's identity trajectory in the direction of the average or beyond the average. In this way, we manipulated a face's values on only the shape or reflectance features, respectively.

Procedure

To assess the contribution of shape and surface reflectance to the overall opponent-identity effect, we measured the adaptation effect with shape- and reflectance-varying faces. In this experiment, we tested only within-viewpoint adaptation. The procedure for this study was similar to that used in Experiment 1, with the following changes. In the shape-varying condition, we replaced the stimuli of the previous experiment, including those used in the no-adaptation condition, with frontal views of shape-varying faces. Again, the antiface was the adapting stimulus, and the anticaricatures and average were the test faces. In the reflectance-varying condition, all adaptation and test images were frontal views of reflectance-varying faces.

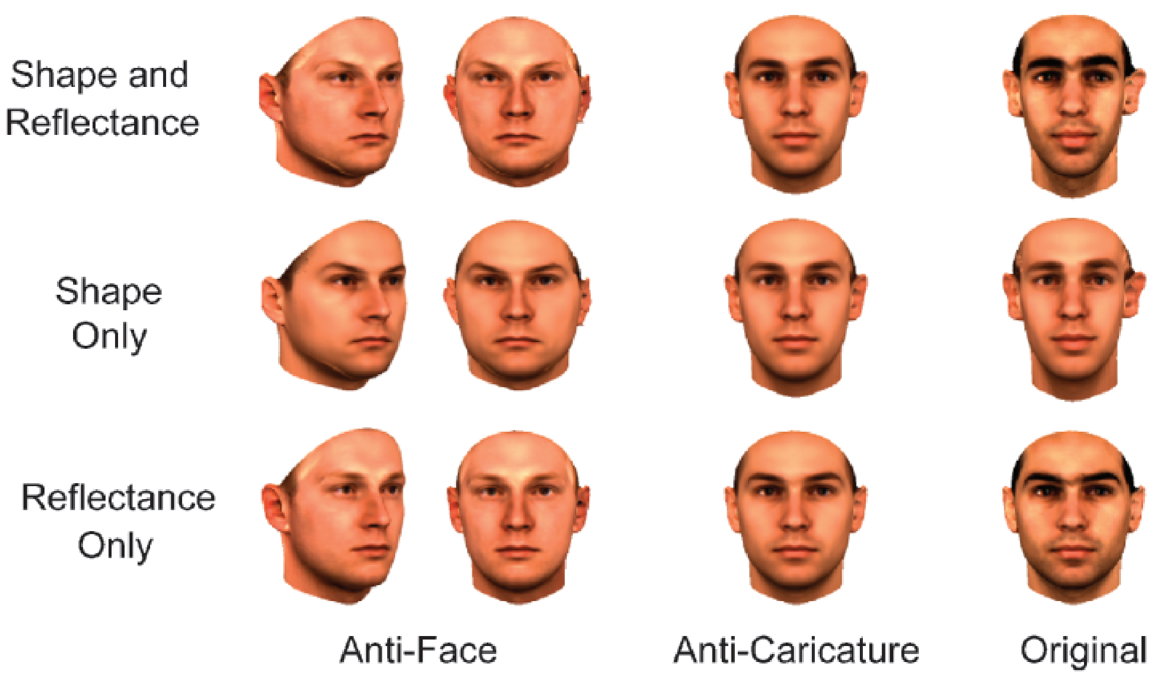

Fig. 3. Manipulation of shape and reflectance information in the face space. The face space encodes both three-dimensional shape and surface reflectance in terms of the positions and the red, green, and blue reflectance values of 75,000 surface points. Shape and reflectance can be varied simultaneously (top row) or separately: Shape-varying (shape-only) stimuli all share the same average reflectance (second row), and reflectance-varying (reflectance-only) stimuli all have the average shape (third row).

Results and Discussion

Adaptation effects were apparent for both the shape- and the reflectance-varying conditions (Fig. 4, within-viewpoint condition). In match trials, anticaricature identification improved with adaptation for both shape-varying faces, $F(1, 18) = 14.73$, $p_{rep} = .986$, $\eta^2 = .18$, and reflectance-varying faces, $F(1, 18) = 15.01$, $p_{rep} = .99$, $\eta^2 = .20$. The average face was identified as the match to the shape-varying adapting stimulus on .52 of the trials, $t(9) = 8.06$, $p_{rep} = .996$, $d = 2.55$. The average face was identified as the match to the reflectance-varying adapting stimulus on .40 of the trials, $t(9) = 3.93$, $p_{rep} = .97$, $d = 1.24$.

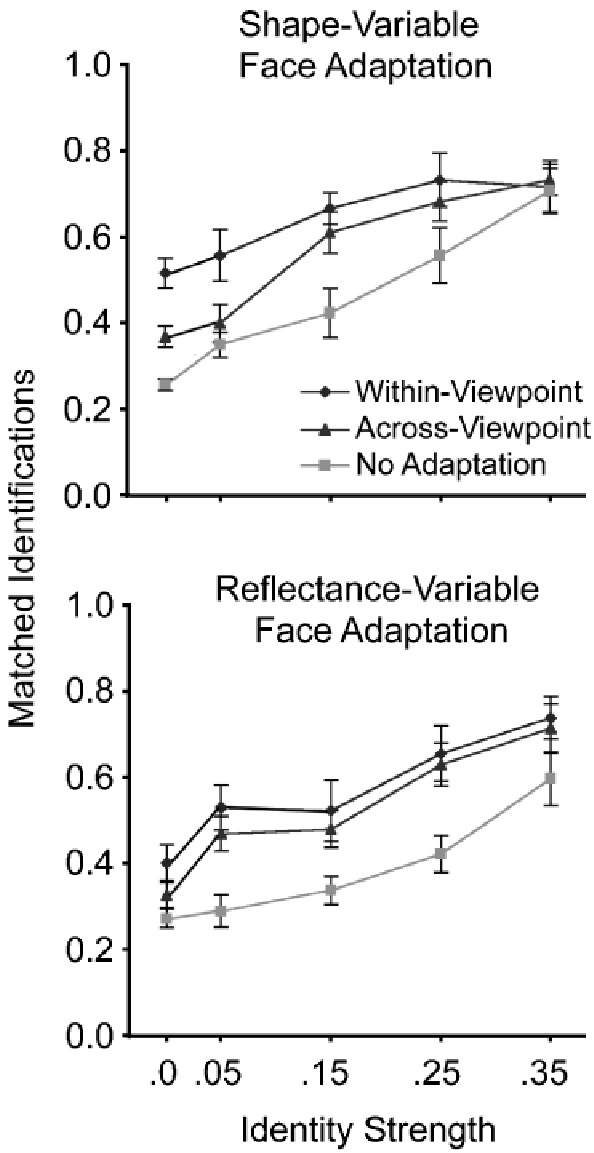

Fig. 4. Results from Experiments 2 and 3: the proportion of trials in the within- and across-viewpoint conditions on which the test face was identified as the match to the adapting antiface stimulus, as a function of the identity strength of the test face. The no-adaptation line indicates the proportion of correct identifications for the faces as a function of identity strength (excluding the 0-identity-strength point, which reflects guessing responses to the average). The top graph presents results for the shape-varying faces, and the bottom graph presents results for the reflectance-varying faces.

The representation of faces that is affected by opponent-identity adaptation, therefore, includes information about both three-dimensional shape and surface reflectance. From the flip side, these findings illustrate that robust and consistent face-adaptation effects can occur in the absence of either identity-specific reflectance information or identity-specific shape information. The fact that neither reflectance nor shape is critical for achieving adaptation clarifies our interpretation of the comparable effectiveness of warped and view-changed faces as adapting stimuli. Specifically, it rules out reflectance overlap as a complete explanation of the view-changed adaptation effect. Indeed, perceptual adaptation to the shape-varying faces shows that it is possible to selectively adapt facial identity when no diagnostic reflectance information is available for discriminating among individuals.

Given that face adaptation involves view-transferable representations, we next considered how shape information and reflectance information contribute to face recognition across changes in viewpoint.

EXPERIMENT 3

We assessed the relative contributions of shape and reflectance information to viewpoint transfer of adaptation by repeating the previous experiment with one change: The adaptation stimuli were view-changed images of the shape-varying and reflectance-varying faces. As in the previous experiment, we tested with frontal anticaricatures of the shape-varying and reflectance-varying faces. Adaptation from a different viewpoint facilitated the identification of both shape-varying faces, $F(1, 23) = 6.69$, $p_{rep} = .93$, $\eta^2 = .05$, and reflectance-varying faces, $F(1, 23) = 10.58$, $p_{rep} = .974$, $\eta^2 = .11$ (see Fig. 4, across-viewpoint condition). Identification of the average face as the match to the adapting face was greater than chance in both the shape-varying condition ($M = .37$), $t(14) = 5.10$, $p_{rep} = .99$, $d = 1.32$, and the reflectance-varying condition ($M = .33$), $t(14) = 2.72$, $p_{rep} = .93$, $d = 0.70$.

These results indicate that both shape and reflectance information can carry information about the identity of a face across a change in three-dimensional viewpoint.

GENERAL DISCUSSION

The novel findings of this study are as follows. First, both the shape and the reflectance information in a face contribute to the identity-adaptation effect, effectively altering the perceptual appearance of subsequent faces viewed from the same viewpoint or from a different viewpoint. This indicates that face representations include both shape and reflectance information in a form that generalizes across changes in three-dimensional viewpoint. This viewpoint generalization occurs, however, via processes that are not entirely specific to faces, but rather may draw on basic perceptual entities composed of features and configurations. The perception of face identity is altered following adaptation to a nonface stimulus that preserves critical perceptual features of the face. This speaks in favor of high-level visual representations that may be shared among objects and faces (Haxby et al., 2001), a conclusion that stands in contrast to the idea that face representations comprise face-specific neural mechanisms that are accessible only to faces. The present data are consistent with a cascaded representation of faces that pools lower-level patterns of shape and reflectance into configural combinations that specify the appearance of a face over a range of viewpoints.

Second, adaptation effects occur for faces lying on opposing sides of an identity trajectory, in either the shape or the reflectance subspace. This indicates that opponent-based processes are at the core of high-level visual codes for faces.

On the basis of these findings, we propose a novel theory of the visual information retained in high-level face representations. This theory extends an opponent-based, contrastive face code to the level of view-transferable processes. We discuss the two components of this theory in terms of their implications for computational models of face recognition and for understanding neural theories of face and object recognition.

First, the finding that surface reflectance is encoded in a way that transfers across three-dimensional viewpoint has implications for the role of reflectance information in neural and computational models of face and object recognition. The present data suggest a more explicit coding of facial reflectance information than has been implemented in most neurally inspired computational models of face recognition (e.g., Biederman, 1987; Marr, 1982; Poggio & Edelman, 1990; Riesenhuber & Poggio, 1999, 2002). The explicit coding of shape has been emphasized as necessary for constancy in predicting the retinal images that result with changes in viewpoint and illumination conditions. Although equally critical, surface-reflectance information is rarely retained following the early visual processing stages of most computational and neural models of object and face recognition. Indeed, human recognition data suggest that reflectance and shape contribute in roughly equal measures to face recognition over changes in viewpoint (O'Toole, Vetter, & Blanz, 1999). It is reasonable, therefore, that both kinds of information be retained at view-transferable levels of encoding.

Second, opponent-based selective adaptation to reflectance and shape information in faces, both within and across viewpoint, suggests that a face representation with underlying opponent processing mechanisms is maintained along the visual hierarchy up to a level of encoding that can operate across a change in viewpoint. We have shown that viewing a particular face from a particular viewpoint can affect the subsequent perception of an entirely different face from a novel viewpoint. This result is consistent with work by Jeffery, Rhodes, and Busey (2006) that shows significant, but diminished adaptation transfer to distorted faces (e.g., Webster & MacLin, 1999) rotated away from the adapting view. Recent work has demonstrated also that the transfer of adaptation across three-dimensional viewpoint increases as a person becomes more familiar with a face (Jiang, Blanz, & O'Toole, in press). The strengthening of transfer is indicative of a representation that evolves with experience. The degree of adaptation transfer, therefore, may speak more to the richness and interconnected nature of the visual representations for faces than to the existence of view-specific versus view-independent neural codes (Booth & Rolls, 1998).

The effect of adaptation on the perception of individual face identities is consistent with a population-based coding of faces. The neural and computational implications of this finding are that a contrastive, relative code, rather than an absolute code, may characterize high-level face representations, up to the level of view-independent responses. The combined adaptation to reflectance and shape in faces may have lower-order visual analogues in well-known form-contingent aftereffects like the McCollough effect (McCollough, 1965)—an opponent-based cross-coupling of reflectance and line orientation that produces a color aftereffect in a neutral gray grating stimulus. As in the case of lower-order visual adaptation, the relational nature of opponent-based coding for faces has advantages in optimizing neural resources by adjusting gain control, normalizing the system to expected input, and decorrelating neural responses to reduce redundancy (Barlow, 1990; Webster, 2003; Webster, Werner, & Field, 2005).

In summary, opponent identity adaptation to computer-generated three-dimensional faces provides a unique and powerful tool for testing predictions about the nature of high-level face representations. Face adaptation transcends viewpoint-specific processes, making it ideal as a method for unraveling the mechanisms of perceptual constancy for faces.

Acknowledgments—Thanks are due to James Cutting and to an anonymous reviewer for suggesting the warped-face control condition in Experiment 1. We also thank Michael Webster and Colin Clifford for helpful comments on a previous version of this manuscript.

REFERENCES

Aguirre, G.K., Singh, R., & D'Esposito, M. (1999). Stimulus inversion and the responses of face and object-sensitive cortical areas. *NeuroReport*, *10*, 189–194.

Barlow, H.B. (1990). A theory about the functional role and synaptic mechanism of visual aftereffects. In C. Blakemore (Ed.), *Vision: Coding and efficiency* (pp. 363–375). Cambridge, England: Cambridge University Press.

Biederman, I. (1987). Recognition by components: A theory of image understanding. *Psychological Review*, *94*, 115–147.

Blanz, V., O'Toole, A.J., Vetter, T., & Wild, H.A. (2000). On the other side of the mean: The perception of dissimilarity in human faces. *Perception*, *29*, 885–891.

Blanz, V., & Vetter, T. (1999, August). *A morphable model for the synthesis of 3D faces*. Paper presented at the annual SIGGRAPH conference, Association for Computing Machinery, Los Angeles.

Booth, M.C.A., & Rolls, E.T. (1998). View-invariant representations of familiar objects by neurons in the inferior temporal visual cortex. *Cerebral Cortex*, *8*, 510–523.

Golby, A.J., Gabrieli, J.D.E., Chiao, J.Y., & Eberhardt, J. (2001). Differential responses in the fusiform region to same-race and other-race faces. *Nature Neuroscience*, *4*, 845–850.

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.

Haxby, J.V., Ungerleider, L.G., Clark, V.P., Schouten, J.L., Hoffman, E.A., & Martin, A. (1999). The effect of face inversion on activity in the human neural systems for face and object perception. *Neuron*, *22*, 188–189.

Hietanen, J.K., Perrett, D.I., Oram, M.W., Benson, P.J., & Dittrich, W.H. (1992). The effects of lighting conditions on responses of cells selective for face views in the macaque temporal cortex. *Experimental Brain Research*, *89*, 157–171.

Jeffery, L., Rhodes, G., & Busey, T. (2006). View-specific coding of face shape. *Psychological Science*, *17*, 501–505.

Jiang, F., Blanz, V., & O'Toole, A.J. (in press). The role of familiarity in view transferability of face identity adaptation. *Journal of Vision*.

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.

Kanwisher, N., Tong, F., & Nakayama, K. (1998). The effect of face inversion on the human fusiform face area. *Cognition*, *68*, B1–B11.

Kourtzi, Z., Erb, M., Grodd, W., & Bülthoff, H.H. (2003). Representation of the perceived 3-D object shape in the human lateral occipital complex. *Cerebral Cortex*, *13*, 911–920.

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.

Logothetis, N.K., Pauls, J., & Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. *Current Biology*, *5*, 552–563.

Marr, D. (1982). *Vision*. San Francisco: Freeman.

McCollough, C. (1965). Color adaptation of edge-detectors in the human visual system. *Science*, *149*, 1115–1116.

Moradi, F., Koch, C., & Shimojo, S. (2005). Face adaptation depends on seeing the face. *Neuron*, *45*, 169–175.

O'Toole, A.J., Vetter, T., & Blanz, V. (1999). Two-dimensional reflectance and three-dimensional shape contributions to recognition of faces across viewpoint. *Vision Research*, *39*, 3145–3155.

Phillips, P.J., Moon, H., Rizvi, S., & Rauss, P. (2000). The FERET evaluation methodology for face recognition algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *22*, 1090–1103.

Poggio, T., & Edelman, S. (1990). A network that learns to recognize 3D objects. *Nature*, *343*, 263–266.

Rhodes, G., Jeffery, L., Watson, T.L., Clifford, C.W.G., & Nakayama, K. (2003). Fitting the mind to the world: Face adaptation and attractiveness aftereffects. *Psychological Science*, *14*, 558–566.

Rhodes, G., Jeffery, L., Watson, T.L., Jaquet, E., Winkler, C., & Clifford, C.W.G. (2004). Orientation-contingent face aftereffects and implications for face-coding mechanisms. *Current Biology*, *14*, 2119–2123.

Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, *2*, 1019–1025.

Riesenhuber, M., & Poggio, T. (2002). Neural mechanisms of object recognition. *Current Opinion in Neurobiology*, *12*, 162–168.

Suzuki, S., & Cavanagh, P. (1998). A shape-contrast effect for briefly presented stimuli. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1–27.

Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Quarterly Journal of Experimental Psychology*, *43A*, 161–204.

Vogels, R., & Biederman, I. (2002). Effects of illumination intensity and direction on object coding in macaque inferior temporal cortex. *Cerebral Cortex*, *12*, 756–766.

Vuilleumier, P., Henson, R.N., Driver, J., & Dolan, R.J. (2002). Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nature Neuroscience*, *5*, 491–499.

Watson, T.L., & Clifford, C.W.G. (2003). Pulling faces: An investigation of the face-distortion aftereffect. *Perception*, *32*, 1109–1116.

Webster, M.A. (2003). Pattern-selective adaptation in color and form perception. In L.M. Chalupa & J.S. Werner (Eds.), *The visual neurosciences* (Vol. 2, pp. 936–947). Cambridge, MA: MIT Press.

Webster, M.A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, *428*, 557–561.

Webster, M.A., & MacLin, O.H. (1999). Figural aftereffects in the perception of faces. *Psychonomic Bulletin & Review*, *6*, 647–653.

Webster, M.A., Werner, J.S., & Field, D.J. (2005). Adaptation and the phenomenology of perception. In C.W.G. Clifford & G. Rhodes (Eds.), *Fitting the mind to the world: Adaptation and aftereffects in high level vision* (pp. 241–277). New York: Oxford University Press.

Young, M.P., & Yamani, S. (1992). Sparse population coding of faces in inferotemporal cortex. *Science*, *256*, 1327–1331.

Zhao, L., & Chubb, C.F. (2001). The size-tuning of the face-distortion aftereffect. *Vision Research*, *41*, 2979–2994.

(Received 4/4/05; Revision accepted 8/1/05;
Final materials received 10/5/05)