PSYCHOLOGICAL SCIENCE

A JOURNAL OF THE ASSOCIATION FOR PSYCHOLOGICAL SCIENCE
aps

Research Article

Three-Dimensional Information in Face Representations Revealed by Identity Aftereffects

Fang Jiang,[1] **Volker Blanz,**[2] **and Alice J. O'Toole**[3]

[1]Unité Cognition et Développement, Université Catholique de Louvain; [2]Fachgruppe Medieninformatik, The University of Siegen; and [3]School of Behavioral and Brain Sciences, The University of Texas at Dallas

ABSTRACT—*Representations of individual faces evolve with experience to support progressively more robust recognition. Knowledge of three-dimensional face structure is required to predict an image of a face as illumination and viewpoint change. Robust recognition across such transformations can be achieved with representations based on multiple two-dimensional views, three-dimensional structure, or both. We used face-identity adaptation in a familiarization paradigm to address a long-standing controversy concerning the role of two-dimensional versus three-dimensional information in face representations. We reasoned that if three-dimensional information is coded in the representations of familiar faces, then learning a new face using images generated by one three-dimensional transformation should enhance the robustness of the representation to another type of three-dimensional transformation. Familiarization with multiple views of faces enhanced the transfer of face-identity adaptation effects across changes in illumination by compensating for a generalization cost at a novel test viewpoint. This finding demonstrates a role for three-dimensional information in representations of familiar faces.*

For an individual face to be recognized, visual information about the face from a two-dimensional projection on the retina must be encoded in the visual cortex and compared with an internal representation of the face. Humans' ability to recognize familiar faces effortlessly despite dramatic changes in appearance with viewing conditions suggests that the visual cortex represents familiar faces in a way that can support recognition invariance. Robust recognition, however, is not characteristic in the case of unfamiliar faces (Hancock, Bruce, & Burton, 2000). Rather, the extent to which face recognition operates invariantly across viewing conditions differs as a function of the amount and kinds of experience an observer has with a given face.

Although it is clear from studies of face recognition and adaptation (cf. Burton, Bruce, & Hancock, 1999; Carbon & Leder, 2005; Carbon et al., 2007; Jiang, Blanz, & O'Toole, 2007) that familiarity enhances the flexibility and robustness of face recognition, the underlying representations that achieve this invariance remain controversial. Traditionally, two theoretical approaches have framed research on the nature of the representations mediating face and object recognition. Viewer-centered theories (e.g., Poggio & Edelman, 1990) propose that the human visual system represents faces and objects using two-dimensional descriptions of different views; recognition is achieved by matching an incoming view to the stored views (Bülthoff & Edelman, 1992). Object-centered theories (e.g., Biederman, 1987) posit that representations take the form of three-dimensional structure information that is independent of viewpoint.

There are two fundamental challenges in interpreting the results of previous studies that support viewer- and object-centered accounts of recognition. The first is that familiarity has rarely been considered explicitly or manipulated in these studies. The second is that in making inferences about the type of information (two- or three-dimensional) coded in the representations, the studies have relied strongly on whether view-specific or view-invariant patterns of behavioral and neural responses are observed. The role of familiarity in enriching individual face representations weakens the logic linking the viewpoint dependency of neural and behavioral responses to the question of whether representations are based on two- or three-

---

Address correspondence to Fang Jiang, Unité Cognition et Développement, Université Catholique de Louvain, 10 Place du Cardinal Mercier, 1348 Louvain-la-Neuve, Belgium, e-mail: fang.jiang@uclouvain.be.

Psychological Science
RESEARCH, THEORY, & APPLICATION IN PSYCHOLOGY AND RELATED SCIENCES

dimensional information. In other words, without control of face familiarity, predictions about view-dependent versus view-independent performance are unspecified.

From a neural perspective, there is evidence for both view-specific and view-invariant coding mechanisms in visual cortex. Face-responsive neurons have been found in various regions of the monkey temporal visual cortex (e.g., Desimone, Albright, Gross, & Bruce, 1984; De Souza, Eifuku, Tamura, Nishijo, & Ono, 2005; Hasselmo, Rolls, & Baylis, 1989; Perrett et al., 1991). Although most face-responsive cells are selective for views of faces and respond only to a specific view or a limited range of views (e.g., Desimone et al., 1984; Perrett et al., 1991), some neurons respond to faces in a view-invariant way (e.g., Hasselmo et al., 1989; Perrett et al., 1991). Studies of face-responsive neurons, therefore, suggest the coexistence of view-specific and view-invariant neuronal responses to faces.

The role of familiarity in neural codes is clear from long-standing evidence that the selectivity of individual face-responsive neurons changes with experience. Specifically, some face neurons that respond uniformly to novel faces show differentiable responses as the faces become familiar (Rolls, Baylis, Hasselmo, & Nalwa, 1989). The role of experience is further indicated by findings from studies of object-responsive cells in monkey inferior temporal (IT) visual cortex. Logothetis, Pauls, and Poggio (1995) found that following extensive laboratory training with multiple views of computer-generated three-dimensional objects, a small number of IT neurons responded to previously unknown objects regardless of their viewpoint. Booth and Rolls (1998) also found IT neurons with view-invariant responses to real-world objects after monkeys had been exposed to the objects from different views over a period of time.

In the behavioral literature, until recently, findings of viewpoint-dependent recognition performance have also been interpreted as evidence for image-based codes or three-dimensional structural codes (see Peissig & Tarr, 2007, for a review). Stankiewicz (2002) and other investigators have argued that the evidential link between the view-specificity of performance and the two-dimensional versus three-dimensional nature of the information encoded in visual representations is weak. A model that is based on a three-dimensional structural description can produce view-specific recognition performance if the process that matches the internal representation with the input data is limited (cf. Liu, Knill, & Kersten, 1995). Alternatively, view-invariant performance can be achieved if enough two-dimensional view templates are stored in the representation to allow for perfect view approximation. Thus, it is not surprising that both view-dependent (e.g., Bülthoff & Edelman, 1992; Tarr, 1995; Tarr & Pinker, 1989) and view-invariant (e.g., Biederman & Bar, 1999; Biederman & Gerhardstein, 1993) recognition performance have been reported (cf. Hayward, 2003).

Behavioral results are consistent with neural studies in that view-dependent behavioral performance is most commonly found for unfamiliar objects, and view-invariant performance is most often observed for familiar objects (Wallis & Bülthoff, 1999). The degree to which recognition is view dependent, therefore, is likely to be mediated by the familiarity of the objects. The increased robustness of more familiar face representations to changes in viewpoint has also been demonstrated using face-adaptation paradigms. In the identity-adaptation paradigm, prolonged exposure to a face with opposite characteristics (a synthetically created "antiface") biases subsequent identification of the veridical face (Leopold, O'Toole, Vetter, & Blanz, 2001). Identity-adaptation effects are thought to tap high-level visual representations because they show tolerance to two-dimensional affine transformations in the size (e.g., Anderson & Wilson, 2005) and retinal position (e.g., Leopold et al., 2001) of the adapting and test stimulus. Identity adaptation also shows partial tolerance to changes in three-dimensional viewpoint between the adapting and test faces (Jiang, Blanz, & O'Toole, 2006). Using identity adaptation with a familiarity manipulation, we (Jiang et al., 2007) found that tolerance to changes in view increased for individual faces as participants became more familiar with the faces. Specifically, familiarity with individual faces enhanced both the overall magnitude of identity aftereffects and the transfer of identity aftereffects across changes in three-dimensional viewpoint. These results suggest that the progression from a view-constrained to a more view-transferable face representation over the course of learning can be measured using an identity-adaptation paradigm.

In sum, when familiarity is not taken into account, the view dependence of performance is not a useful index of the nature of the information coded in neural and perceptual representations of faces. In the study reported here, we used identity adaptation as a novel, direct approach to assessing the encoding of three-dimensional information in individual face representations as they evolve through experience. We reasoned that if three-dimensional information is coded in representations of familiar faces, then learning a new face using images generated by one transformation referencing three-dimensional structure should enhance the robustness of the representation to another, unrelated transformation referencing three-dimensional structure. Critical to the logic of our study is that changes in the appearance of a face induced by viewpoint variation and changes in the appearance of a face induced by illumination variation both depend directly on the three-dimensional structure of the face. Although we do not assume that viewpoint and illumination transformations are qualitatively equivalent in all ways (cf. Discussion), their common dependence on three-dimensional structure supports the following prediction. If knowledge about three-dimensional face structure is acquired from exposure to multiple views, such exposure should support the transfer of identity adaptation over changes in illumination. This would provide evidence for the coding of three-dimensional face structure in representations of familiar faces.

We report two experiments. Experiment 1 demonstrates that adaptation transfers partially across illumination change, even with minimally familiar faces. This prerequisite study was needed to establish parity between viewpoint and illumination. The second, more interesting experiment addresses the nature of the information in face representations. The design for the second experiment involved a manipulation of face familiarity in terms of both the number of exposures and the viewpoint of the training faces in the learning session and a manipulation of the consistency of illumination direction in a face-adaptation test. First, participants were familiarized with faces shown either from a single view (i.e., the adaptation-test view) or from multiple views (i.e., two views outside the view used in the adaptation test; see Table 1). We also varied the number of exposures between two single-view conditions to test for changes in the magnitude of adaptation and adaptation transfer with familiarity. Next, we measured adaptation effects for the test view both when the illumination of the test face was consistent with the illumination of the adapting face and when the illumination of the test face was inconsistent with the illumination of the adapting face.

GENERAL METHOD

Participants

Eighty-five undergraduate students at The University of Texas at Dallas participated (32 in Experiment 1 and 53 in Experiment 2) after providing written consent in accordance with procedures approved by the institutional review board. Data from 11 participants were deleted because of uniformly low performance. Data from another 2 participants were incomplete because of a computer problem. The final samples included 30 participants in Experiment 1 (10 in each adaptation condition) and 42 participants in Experiment 2 (14 in each familiarity condition).

Stimuli

Face stimuli were created using a three-dimensional morphable model (Blanz & Vetter, 1999). Built on the concept of a multidimensional prototype-centered face space, this model codes each face by its direction and distance from the average face (cf. Valentine, 1991). The direction of a face vector defines its identity, and its distance from the average face determines its identity strength (i.e., distinctiveness). This prototype-based theory of face representation has been described elsewhere (e.g., Jiang et al., 2006; Leopold et al., 2001) and has been supported by both functional magnetic resonance imaging (e.g., Loffler, Yourganov, Wilkinson, & Wilson, 2005) and single-unit neuronal (Leopold, Bondar, & Giese, 2006) evidence.

TABLE 1

Familiarity Conditions in Experiment 2

Condition	Description
Single-view, 4-exposures	4 exposures of a 22.5°-rotated view
Single-view, 16-exposures	16 exposures of a 22.5°-rotated view
Multiple-views, 16-exposures	8 exposures of a frontal view and 8 exposures of a 45°-rotated view

We used four male faces from our previous studies (Jiang et al., 2006, 2007) as the veridical faces. By morphing the veridical faces (identity strength = 1) toward the average face (identity strength = 0), we generated anticaricatures (i.e., less distinctive versions of the veridical faces) with identity strength varying from .05 to .35 in .10 steps. Antifaces (i.e., opposites of the veridical faces) with negative identity strength (−.75) were created by morphing the veridical faces further beyond the average face. Figure 1a presents some examples of the stimuli.

The three-dimensional model allowed us to create images using different illumination and viewpoint conditions. Illuminations were modeled using computer graphics by setting the direct light from the top left or the top right of the camera, keeping the overall luminance well controlled. Viewpoints were manipulated by rotating the face stimuli in depth.

Antifaces were used as adapting faces, which were illuminated either from the top left or from the top right. The average face and anticaricatures were used as test faces, and they were always illuminated from the top left. Because of the nonfrontal pose (22.5°) in the test trials, top-left and top-right lighting differed by more than a simple mirror reflection. We used ve-

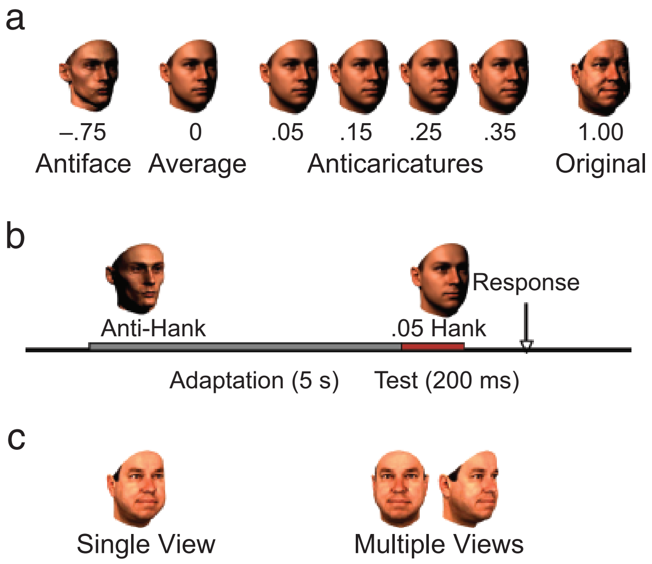

Fig. 1. Stimuli and paradigm: (a) example of the face stimuli created along a particular identity trajectory, (b) the identity-adaptation paradigm (modified from Leopold, O'Toole, Vetter, & Blanz, 2001), and (c) examples of the training faces used in Experiment 2. In (a), the numbers below the images indicate identity strength. Morphing the veridical faces (identity strength = 1) toward the average face (identity strength = 0) generated anticaricatures (i.e., less distinctive versions of the veridical faces). Antifaces (i.e., opposites of the veridical faces) with negative identity strength (−.75) were created by morphing the veridical faces further beyond the average face. In the inconsistent-illumination trial shown in (b), the adapting face is illuminated from the top right, and the test face is illuminated from the top left. Following 5 s of adaptation to an antiface, participants were asked to identify a briefly presented (200 ms) test face as one of the four learned veridical faces.

ridical faces and their .35 anticaricatures as the training faces. All training faces in Experiment 1 were illuminated from the top left. For the training faces used in Experiment 2, ambient lights were applied (Fig. 1c) to enhance the effect of familiarization.

General Procedure

Each experiment began with a short training session, followed by practice and test sessions. The training proceeded as follows. First, participants were asked to learn names for four veridical faces. These four faces were presented with their names a number of times (varied between experiments) for 5 s each in random order. Then, with feedback, participants named each original face once by pressing a labeled key on the keyboard. This procedure was repeated once with .35 anticaricatures. Immediately after training, a practice session was provided to familiarize participants with the task. Practice trials had the same format as test trials, but only .35 anticaricatures were used. These trials were excluded from the analysis.

In the test session, test faces were flashed briefly (200 ms), and participants identified each test face as one of the four veridical faces (Fig. 1b). All test stimuli were illuminated from the top left. Each test face was presented following 5 s of a blank screen (*no adaptation*), 5 s of adaptation to an antiface illuminated from the top left (*consistent-illumination adaptation*), or 5 s of adaptation to an antiface illuminated from the top right (*inconsistent-illumination adaptation*). Both the adapting and the test faces were rotated 22.5° to the right. An example of an inconsistent-illumination adaptation trial appears in Figure 1b.

EXPERIMENT 1

Our goal in Experiment 1 was to determine whether face-identity adaptation effects transfer partially across changes in illumination for faces with which observers do not have extensive familiarity. This would make illumination comparable to viewpoint in showing some transfer with minimally familiar faces (Jiang et al., 2006). We measured identity aftereffects as a function of identity strength both within and across illumination conditions.

Method

During training, the number of exposures and views of the training faces were fixed (four exposures to the 22.5° view for each training face) for all the participants. The training session was comparable to that used in our previous viewpoint experiment (Jiang et al., 2006) and was necessary for participants to perform the identification task (i.e., naming the faces). Immediately after training, 32 practice trials were given.

In the test session, adaptation effects were measured as a function of identity strength of the test face, which varied within participants. Five identity strengths were tested: 0 (i.e., the average face), .05, .15, .25, and .35 (see Fig. 1a). Adaptation condition varied across participants. For participants in the no-adaptation condition, equal numbers of neutral trials (i.e., the test face was the identity-neutral average face) and anticaricature trials (i.e., the test face was an anticaricature) were presented, for a total of 96 trials. For participants in the adaptation conditions, equal numbers of neutral trials, match trials (i.e., the adapting and test faces were from the same identity), and nonmatch trials (i.e., the adapting and test faces were from two different identities) were presented, for a total of 144 trials. Only the data from the match and neutral trials were included in the analysis. Trial order was randomized for each participant.

Results

We used the proportion of matched identifications (i.e., the proportion of trials in which the test face was identified as the match to the adapting face) to measure identity adaptation and adaptation transfer, for comparability with our previous study (Jiang et al., 2006). As expected, for identity strengths less than .15, adaptation facilitated identification when the adapting and test faces were illuminated from the same direction, $F(1, 18) = 23.19$, $p_{\text{rep}} = .986$, $\eta^2 = .416$, and also when the two faces were illuminated from different directions, $F(1, 18) = 11.34$, $p_{\text{rep}} = .974$, $\eta^2 = .238$ (see Fig. 2). For the same range of identity strengths, the magnitude of adaptation effects was diminished in the inconsistent-illumination condition compared with the consistent-illumination condition, $F(1, 18) = 3.54$, $p_{\text{rep}} = .844$, $\eta^2 = .116$. Thus, we found that identity adaptation transfers partially across changes in illumination, just as it transfers

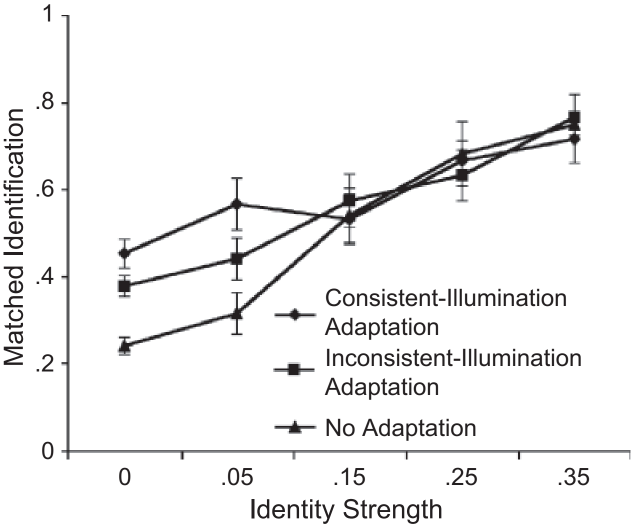

Fig. 2. Results from Experiment 1. For the consistent- and inconsistent-illumination adaptation conditions, the graph shows the proportion of trials in which the test face was identified as the match to the adapting antiface, as a function of the identity strength of the test face. For the no-adaptation condition, the graph shows the proportion of correct identifications of the test faces, as a function of their identity strength. Note that for 0 identity strength, performance in the no-adaptation condition reflected chance performance, as the average face had no identity. Error bars indicate standard errors.

partially across changes in three-dimensional viewpoint (Jiang et al., 2006).

EXPERIMENT 2

The purpose of this experiment was to test whether learning a new face using images generated by one three-dimensional transformation (viewpoint transformation) would enhance adaptation transfer over another type of three-dimensional transformation (illumination transformation). To investigate whether experience with multiple views enhances adaptation transfer across changes in illumination, we manipulated the way participants learned the faces.

Method

Participants were assigned to one of three familiarity conditions that varied both in the number of exposures to each face and in the viewpoints of the faces during training (see Table 1). In the *single-view, 4-exposures* condition, participants saw each 22.5°-rotated training face four times. In the *single-view, 16-exposures* condition, participants viewed each 22.5°-rotated training face 16 times. In the *multiple-views, 16-exposures* condition, each training face was shown 16 times, 8 times with a frontal view and 8 times with a 45°-rotated view. Note that it was only in the single-view familiarity conditions that training views were the same as the adaptation-test views. Immediately following the training session, 40 practice trials were given.

Next, we tested identity-adaptation effects as a function of the consistency of the illumination between the adapting and test faces. The magnitude of adaptation effects was assessed for test faces in the 22.5° view at a single low identity strength (.05). We chose this identity strength because there were clear adaptation and adaptation-transfer effects at this level in Experiment 1. Regardless of familiarity condition, 144 trials consisting of equal numbers of no-adaptation, consistent-illumination adaptation, and inconsistent-illumination adaptation trials were presented to each participant. In order to avoid effects of adaptation to low-level light (induced by switching between adapting faces lit from the top right and lit from the top left), we presented trials from different adaptation conditions in separate blocks. The no-adaptation block was always presented at the end. The order of the consistent- and inconsistent-illumination blocks was counterbalanced across participants. Equal numbers of match and nonmatch trials were included in adaptation blocks; nonmatch trials were excluded from analysis.

Results

For each participant, the strength of identity adaptation was measured by computing the difference between identification performance after adaptation and identification performance in the no-adaptation condition. Two such adaptation scores were calculated, one for each adaptation condition (cf. Jiang et al., 2007). We used this measure to control for possible increases in sensitivity to individual faces as they became more familiar. This control was needed because our predictions concerned the strength of adaptation and adaptation transfer across face representations, rather than absolute identification accuracy. Using each participant's own no-adaptation performance as a baseline removed sensitivity increases, leaving a clearer measure of the magnitude of adaptation effects.

The design of Experiment 2 allowed us to address three questions. The first was whether familiarity with a single view enhances identity adaptation and its transfer across a change in illumination. This would be indicated by a stronger adaptation effect as the total number of exposures to a face in the test view increased from 4 to 16. As illustrated in Figure 3, adaptation scores increased significantly from the single-view, 4-exposures condition to the single-view, 16-exposures condition, both in the consistent-illumination adaptation condition, $t(26) = 3.13$, $p_{\text{rep}} = .968$, $d = 1.18$, and in the inconsistent-illumination adaptation condition, $t(26) = 2.52$, $p_{\text{rep}} = .931$, $d = 0.95$. The magnitude of identity aftereffects and the transfer of adaptation increased following familiarization with a single image of a face.

The second question was whether there is a generalization cost for adaptation effects tested at a novel view situated between two learned views. Specifically, both the adaptation and the test faces were presented at the 22.5° view, which was intermediate to the views learned in the multiple-views, 16-exposures condition. If participants in that condition formed an intermediate, or "interpolated," representation without cost, then the magnitude of the adaptation effect in the consistent-illumination condition would have been comparable for those participants and for participants in the single-view, 16-exposures condition. As Figure 3 shows, our results were consistent with an interpolation cost; there was a trend indicating a de-

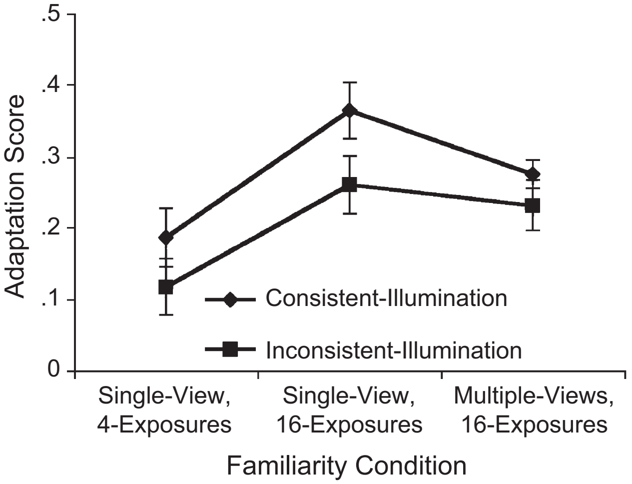

Fig. 3. Effect of face learning on adaptation strength in Experiment 2. Adaptation scores are plotted as a function of familiarity condition, separately for the consistent-illumination and inconsistent-illumination adaptation conditions. Error bars indicate standard errors.

crease in the strength of consistent-illumination adaptation between the single-view, 16-exposures condition and the multiple-views, 16-exposures condition, $t(26) = -2.02, p_{\text{rep}} = .872$, $d = 0.76$. Thus, when the learning consisted of the same number of exposures, but to views situated outside the view tested, the strength of the adaptation effect was diminished.

The central question of this study was whether experience with multiple views enhances adaptation transfer across a change in illumination in a novel test view. If so, the benefit of experience with multiple views would be expected to compensate for the interpolation cost in the novel test view and to boost adaptation transfer across the change in illumination. A generalization cost without compensation would be evident if the adaptation effects in the inconsistent-illumination condition were diminished in the multiple-views, 16-exposures condition relative to the single-view, 16-exposures condition. As Figure 3 shows, however, the strength of the adaptation effect in the inconsistent-illumination condition was comparable in the multiple-views, 16-exposures condition and the single-view, 16-exposures condition, $t(26) = -0.55$, n.s. In fact, the magnitude of this effect was indistinguishable from the magnitude of the adaptation with consistent illumination in the multiple-views, 16-exposures condition, $t(13) = 1.82$, n.s., paired *t* test. This indicates a contribution of illumination-insensitive representations in the human visual system when participants learn faces from multiple views. Therefore, the mechanisms of generalization across viewpoint and generalization across illumination are coupled. The lack of a generalization cost for adaptation transfer across a change in illumination in the novel test view is consistent with a contribution from three-dimensional information learned from multiple views.

DISCUSSION

The results of this study demonstrate that learning three-dimensional face structure from multiple viewpoints can support robust recognition when faces undergo a different transformation that also depends on three-dimensional structure. The evidence for this conclusion is the finding that learning multiple views enhanced the transfer of adaptation effects across illumination changes by compensating for a generalization cost at a novel viewpoint. The use of three-dimensional structure demonstrated here complements a previously established role for two-dimensional view information in face representations (e.g., Fang & He, 2005; Jeffery, Rhodes, Busey, 2006; Wallraven, Schwaninger, Schumacher, & Bülthoff, 2002). The encoding of unfamiliar faces is constrained by the initial views (Bruce & Young, 1986). Encoding of view-based, two-dimensional information about unfamiliar faces, however, does not preclude the encoding of some three-dimensional information to support transformation invariance as a face becomes familiar.

Although the method we employed here makes use of the fact that both viewpoint and illumination variations reference three-dimensional face structure, it is clear that these two transformations are not entirely equivalent. Varying illumination induces complex changes in both shading and intrinsic cast shadows in a face image. These changes can result in larger image differences than the changes that result from varying viewpoint (Tarr, Kersten, & Bülthoff, 1998). Also, changes in illumination have a greater impact on face recognition than on object recognition (cf. Vogels & Biederman, 2002) and can mediate the effects of changes in viewpoint (Hill & Bruce, 1996). Nevertheless, face recognition across one three-dimensional transformation may be supported by face-structure information learned through experience with another three-dimensional transformation.

We also found that adaptation and the transfer of adaptation effects across changes in illumination are strengthened following familiarization with even a single view of a face. This increases our confidence in our previous finding that familiarity with faces increases the magnitude of adaptation and the degree of adaptation transfer across changes in viewpoint (Jiang et al., 2007). We suggested previously that the role of familiarity in increasing the overall magnitude of adaptation effects may reflect enhanced malleability and perceptual accessibility of representations of familiar faces. We also hypothesized that experience with a face from a single view may enhance the ability to use general knowledge about faces for supporting some degree of view transfer (Blanz, Grother, Phillips, & Vetter, 2005). The present finding further supports this hypothesis, suggesting that general experience is essential for coping with variation in appearance due to changes in viewing conditions when initial encounters with a face have been limited. This suggested role of experience in shaping the robustness of face representations is consistent with the findings of Carbon and his colleagues (Carbon & Leder, 2005; Carbon et al., 2007). Using highly familiar faces (i.e., faces of celebrities), they showed that face aftereffects can transfer across different images of the same person, which suggests that internal representations of familiar faces can be adapted or refined constantly to accommodate changing visual inputs.

One surprising finding is that the formation of an intermediate, or interpolated, representation entailed a generalization cost. This was evident from the diminished adaptation effect observed in the consistent-illumination condition when the novel test view was centered between two learned views. Consistent with the two-dimensional view interpolation theory, this finding suggests that extensive exposures to two learning views might be needed for flawlessly interpolating a representation of a face in a novel, intermediate view. The diminished adaptation effects for the unlearned view are also consistent with a three-dimensional model, if one assumes that building a complete three-dimensional representation takes time and experience. The cost we found in Experiment 2 was more noticeable than the relatively minimal interpolation cost found in a previous study with a recognition paradigm (Wallraven et al., 2002). This discrepancy is most likely due to a difference in the paradigms used. When the adapting and test faces are both presented at the intermediate view, the adaptation paradigm directly taps the intermediate representation of a face with a perceptual, rather than memory-based, test.

In conclusion, our findings demonstrate the existence of face representations that become simultaneously robust to two different transformations referencing three-dimensional structure. These findings indicate a role for three-dimensional information in the representations of familiar faces. They further support the unique power of face adaptation (Webster & MacLin, 1999) as a tool for tracking the nature of the visual information encoded in face representations as they evolve with experience.

REFERENCES

Anderson, N.D., & Wilson, H.R. (2005). The nature of synthetic face adaptation. *Vision Research*, *45*, 1815–1828.

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

Biederman, I., & Bar, M. (1999). One-shot viewpoint invariance in matching novel objects. *Vision Research*, *39*, 2885–2899.

Biederman, I., & Gerhardstein, P.C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 1162–1182.

Blanz, V., Grother, P., Phillips, J., & Vetter, T. (2005). Face recognition based on frontal views generated from non-frontal images. In *CVPR'05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (Vol. 2, pp. 454–461). San Diego, CA: IEEE Computer Society.

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.

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.

Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, *77*, 305–327.

Bülthoff, H.H., & Edelman, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Sciences, USA*, *89*, 60–64.

Burton, A.M., Bruce, V., & Hancock, P.J.B. (1999). From pixels to people: A model of familiar face recognition. *Cognitive Science*, *23*, 1–31.

Carbon, C.C., & Leder, H. (2005). Face adaptation: Changing stable representations of familiar faces within minutes? *Advances in Cognitive Psychology*, *1*, 1–7.

Carbon, C.C., Strobach, T., Langton, S.R.H., Harsányi, G., Leder, H., & Kovács, G. (2007). Adaptation effects of highly familiar faces: Immediate and long lasting. *Memory & Cognition*, *35*, 1966–1976.

Desimone, R., Albright, T.D., Gross, C.G., & Bruce, C.J. (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. *Journal of Neuroscience*, *4*, 2051–2068.

De Souza, W.C., Eifuku, S., Tamura, R., Nishijo, H., & Ono, T. (2005). Differential characteristics of face neuron responses within the anterior superior temporal sulcus of macaques. *Journal of Neurophysiology*, *94*, 1252–1266.

Fang, F., & He, S. (2005). Viewer-centered object representation in the human visual system revealed by viewpoint aftereffects. *Neuron*, *45*, 793–800.

Hancock, P.J.B., Bruce, V., & Burton, A.M. (2000). Recognition of unfamiliar faces. *Trends in Cognitive Sciences*, *4*, 330–337.

Hasselmo, M.E., Rolls, E.T., & Baylis, G.C. (1989). The role of expression and identity in the face selective response of neurons in the temporal visual cortex of the monkey. *Behavioral Brain Research*, *32*, 203–218.

Hayward, W.G. (2003). After the viewpoint debate: Where next in object recognition. *Trends in Cognitive Sciences*, *7*, 425–427.

Hill, H., & Bruce, V. (1996). Effects of lighting on the perception of facial surfaces. *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 986–1004.

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. (2006). Probing the visual representation of faces with adaptation: A view from the other side of the mean. *Psychological Science*, *17*, 493–500.

Jiang, F., Blanz, V., & O'Toole, A.J. (2007). The role of familiarity in three-dimensional view transferability of face identity adaptation. *Vision Research*, *47*, 525–531.

Leopold, D.A., Bondar, I.V., & Giese, M.A. (2006). Norm-based face encoding by single neurons in the monkey inferotemporal cortex. *Nature*, *442*, 572–575.

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.

Liu, Z., Knill, D., & Kersten, D. (1995). Object classification for human and ideal observers. *Vision Research*, *35*, 549–568.

Loffler, G., Yourganov, G., Wilkinson, F., & Wilson, H.R. (2005). fMRI evidence for the neural representation of faces. *Nature Neuroscience*, *8*, 1386–1390.

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

Peissig, J.J., & Tarr, M.J. (2007). Visual object recognition: Do we know more now than we did 20 years ago? *Annual Review of Psychology*, *58*, 75–96.

Perrett, D.I., Oran, M.V., Harries, M.H., Bevan, R., Hietanen, J.K., Benson, P.J., & Thomas, S. (1991). Viewer-centred and object-centred coding of heads in the macaque temporal cortex. *Experimental Brain Research*, *86*, 159–173.

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

Rolls, E.T., Baylis, G.C., Hasselmo, M.E., & Nalwa, V. (1989). The effect of learning on the face selective responses of neurons in the cortex in the superior temporal sulcus of the monkey. *Experimental Brain Research*, *76*, 153–164.

Stankiewicz, B.J. (2002). Empirical evidence for independent dimensions in the visual representation of three-dimensional shape. *Journal of Experimental Psychology: Human Perception and Performance*, *28*, 913–932.

Tarr, M.J. (1995). Rotating objects to recognize them: A case study of the role of viewpoint dependency in the recognition of three-dimensional objects. *Psychonomic Bulletin & Review*, *2*, 55–82.

Tarr, M.J., Kersten, D., & Bülthoff, H.H. (1998). Why the visual recognition system might encode the effects of illumination. *Vision Research*, *38*, 2259–2275.

Tarr, M.J., & Pinker, S. (1989). Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology*, *21*, 233–282.

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.

Wallis, G., & Bülthoff, H.H. (1999). Learning to recognize objects. *Trends in Cognitive Sciences*, *3*, 22–31.

Wallraven, C., Schwaninger, A., Schumacher, S., & Bülthoff, H.H. (2002). View-based recognition of faces in man and machine: Revisiting inter-extra-ortho. In H.H. Bülthoff, S.W. Lee, T. Poggio, & C. Wallraven (Eds.), *Proceedings of the Second International Workshop on Biologically Motivated Computer Vision (BMCV 2002)* (pp. 651–660). Berlin, Germany: Springer.

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

(RECEIVED 11/27/07; REVISION ACCEPTED 8/2/08)