ELSEVIER

VISION
RESEARCH

An International Journal for Functional Aspects of Vision

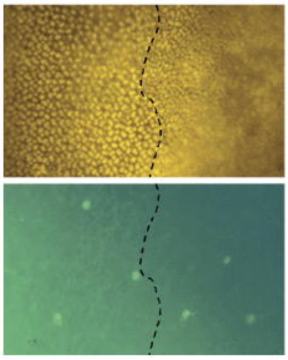

Biochemistry & Cell Biology • Molecular Biology & Genetics
Anatomy, Physiology, Pathology & Pharmacology • Optics, Accommodation & Refractive Error
Circuitry & Pathways • Psychophysics • Perception • Attention & Cognition
Computational Vision • Eye Movements & Visuomotor Control

ISSN 0042-6989 | Volume 47 | Number 4 | February 2007

Available online at www.sciencedirect.com
ScienceDirect
Vision Research 47 (2007) 525–531
Vision Research
www.elsevier.com/locate/visres

The role of familiarity in three-dimensional view-transferability of face identity adaptation

Fang Jiang [a,*], Volker Blanz [b], Alice J. O'Toole [a]

[a] *School of Behavioral and Brain Sciences, GR4.1 The University of Texas at Dallas, Richardson, TX 75083-0688, USA*
[b] *Fachgruppe Medieninformatik Fachbereich, 12 The University of Siegen, 57068 Siegen, Germany*

Received 3 August 2006; received in revised form 11 October 2006

Abstract

Recent studies show that face adaptation effects partially transfer across three-dimensional viewpoint change. Here we investigated whether the degree of adaptation transfer is mediated by experience with a face. We manipulated face familiarity and measured identity aftereffects both within- and across-viewpoint. Familiarity enhanced the overall strength of identity adaptation as well as the degree to which adaptation transferred across-viewpoint change. These findings support the idea that transfer effects in adaptation vary as a function of experience with particular faces, and suggest the use of adaptation as a tool for tracking face representations as they develop.

Keywords: Adaptation; Aftereffect; Face; Familiarity

1. Introduction

Recent studies indicate that adaptation is useful for probing the high-level representations of objects and faces (e.g., Suzuki & Cavanagh, 1998; Webster & MacLin, 1999). Face adaptation effects occur when pre-exposure to a face with particular characteristics biases the perception of a subsequently presented face in a contrastive fashion. For example, the perception of identity can be altered following adaptation to an "anti-face"—a synthetically created opposite of an original face (Leopold, O'Toole, Vetter, & Blanz, 2001). In this case, adaptation to an anti-face facilitates the perception of the original face. Contrastive aftereffects have been found also for face configuration (Webster & MacLin, 1999), natural categories such as gender, race, and expression (Webster, Kaping, Mizokami, & Duhamel, 2004), and for face attractiveness (Rhodes, Jeffery, Watson, Clifford, & Nakayama, 2003).

Compared to traditional aftereffects for basic visual dimensions (e.g., color), aftereffects induced by complex shapes, including faces, are thought to reflect aspects of "high-level" visual processing (Leopold et al., 2001). Convergent evidence for this claim comes from three kinds of findings. First, a conscious perception of the adapting face is essential for identity-specific aftereffects (Moradi, Koch, & Shimojo, 2005). Specifically, Moradi et al. (2005) used a binocular suppression paradigm and showed that identity aftereffects vanished when the adapting face was "invisible" (i.e., not consciously perceived) for more than 3 s. In contrast, low-level orientation aftereffects remain intact under the same suppression (e.g., Moradi et al., 2005).

Second, face aftereffects show tolerance to two-dimensional affine transformations in the size, orientation, and retinal position of a face (Anderson & Wilson, 2005; Jeffery, Rhodes, & Busey, 2006; Leopold et al., 2001; Rhodes et al., 2003; Watson & Clifford, 2003; Zhao & Chubb, 2001). Concomitantly, face processing is known from functional neuroimaging studies to recruit high-level visual areas, such as the fusiform face area (FFA, Kanwisher, McDermott, & Chun, 1997), which do not have a strong retinotopic organization. Responses to faces in the FFA show a substantial degree of tolerance to size and position changes (Grill-Spector et al., 1999; Grill-Spector &

* Corresponding author. Fax: +1 972 883 2491.
E-mail address: fxj018100@utdallas.edu (F. Jiang).

0042-6989/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2006.10.012

Malach, 2001). The transferability of face adaptation effects over two-dimensional affine transformations, therefore, suggests a locus in high-level visual areas.

Third, in addition to surviving two-dimensional affine transformations, partial transfer of face adaptation over changes in three-dimensional viewing parameters also has been reported in recent studies (Jeffery et al., 2006; Jiang, Blanz, & O'Toole, 2006). For example, Jiang et al. (2006) investigated the view-transferability of face codes using the identity-specific adaptation paradigm first introduced by Leopold et al. (2001). In this paradigm, identity aftereffects were measured by having participants adapt to anti-faces and testing their ability to identify faces from anti-caricatures. Adaptation effects are recorded when adaptation to an anti-face facilitates the identification of the anti-caricatures of the original face (Anderson & Wilson, 2005; Jiang et al., 2006; Leopold et al., 2001; Moradi et al., 2005).

In the three-dimensional view-transferability study, Jiang et al. (2006) tested participants in a within-viewpoint adaptation condition and an across-viewpoint adaptation condition. Participants in the within-viewpoint adaptation condition identified frontal anti-caricatures following 5 s of adaptation to frontal views of anti-faces. Participants in the across-viewpoint adaptation condition performed the same task, but with 5 s of adaptation to anti-faces that were rotated 30° to the right. Jiang et al. (2006) found that identity adaptation transferred across a 30° change in viewpoint, though the aftereffects were diminished by comparison to the within-view condition.

Jiang et al. (2006) further assessed the relative contributions of shape and surface reflectance to the identity adaptation effect and to its transfer across-viewpoint change. This was done using face morphs that selectively varied in three-dimensional shape or surface reflectance (Blanz & Vetter, 1999). Specifically, shape-varying faces had their original shapes, combined with the reflectance of the average face. Reflectance-varying faces had their original reflectance combined with the average shape. Jiang et al. (2006) found that both shape and surface reflectance contributed substantially to the identity aftereffect. Moreover, selective adaptation to both shape and reflectance information transferred across changes in three-dimensional viewpoint, indicating that they both have critical roles in a perceptually constant representation of objects and faces.

A similar finding of three-dimensional transfer was reported by Jeffery et al. (2006) using figural shape aftereffects (e.g.,Webster & MacLin, 1999). They found significant, but reduced aftereffects, when the adapting and test views differed by 45°. The adaptation transfer remained diminished when the adapting and test faces were mirrored (−45° to 45°).

The finding of partial transfer of adaptation across three-dimensional viewpoint further supports the idea that the source of face adaptation can be located at relatively high-level visual areas. Inferotemporal cortex (IT) has been considered a plausible neural locus for the neural representations affected by adaptation, because neurons there have face-selective and shift-invariant responses (Leopold et al., 2001). Concomitantly, IT neurons exhibit partial invariance to three-dimensional transformations of objects, as revealed by neurophysiological studies (Booth & Rolls, 1998; Logothetis, Pauls, Bülthoff, & Poggio, 1994, 1995).

Combined, these lines of evidence offer support for the claim that face adaptation taps relatively high-level visual processing mechanisms. The localization of face adaptation effects in high-level visual areas, however, does not guarantee a completely convergent interpretation of results on three-dimensional viewpoint transfer. Whereas Jeffery et al. (2006) argue that partial transfer of adaptation effects is evidence for view-specificity in face representations, Jiang et al. (2006) suggest that the degree of transfer may reflect a changeable property of face representations. In particular, one important factor that has not been considered in these two previous studies is the role of familiarity in the degree of adaptation transfer. It is well known from psychophysical studies that experience with faces, ultimately, allows us to recognize them regardless of viewing condition (see Burton, Bruce, & Hancock, 1999 for a review). Moreover, as revealed by fMRI priming studies, view-invariant coding of faces in human visual cortex is enhanced by familiarity (Eger, Schweinberger, Dolan, & Henson, 2005; Pourtois, Schwartz, Seghier, Lazeyras, & Vuilleumier, 2005).

Neurophysiological findings from single cell studies provide further support for the role of familiarity in the degree of invariance exhibited by neurons in macaque IT cortex. For example, Booth and Rolls (1998) found a small population of IT neurons that responded similarly to different views of familiar objects, which macaques had played with for weeks prior to neurophysiological experiments. View-invariant responses can be derived also from extensive training with a limited number of views, possibly as the outcome of the convergent responses of view-selective neurons (Logothetis et al., 1994, Logothetis, Pauls, & Poggio, 1995). The evolution of a view-invariant code from view-selective inputs is consistent also with hierarchical models of object recognition, such as the one proposed by Riesenhuber and Poggio (1999, 2000). In their model, a view-invariant representation can be formed from a hierarchy of view-tuned mechanisms, via learning to associate view-specific two-dimensional view-specific representations.

The key to interpreting partial view transfer effects in adaptation, therefore, may lie in understanding how familiarity or experience with faces enriches their representations. The degree to which face codes operate invariantly, and consequently the degree to which adaptation transfers across-viewpoint, may change continuously as we gain experience with faces. In the present study, we used identity adaptation to probe the 3D generalizability of face representations as they evolve with experience. Specifically, we investigated the effects of familiarity on the three-dimensional view-transferability of identity adaptation. We hypothesized that if the three-dimensional view-transferability of identity adaptation is modulated by familiarity, the degree of three-dimensional transfer that occurs in face adaptation studies should increase with experience. To address this issue, we manipulated participants' experience with faces and measured identity adaptation effects within and across-viewpoint conditions.

2. Materials and methods

2.1. Materials

Face stimuli were generated using a three-dimensional morphable model developed by Blanz and Vetter (1999). This model is implemented on a multi-dimensional prototype-centered face space in which the identity of an individual is coded by a particular trajectory and the distinctiveness of a face is defined by its distance to the average face (cf., Anderson & Wilson, 2005; Leopold et al., 2001; Jiang et al., 2006; Valentine, 1991; cf. also Giese & Leopold, 2005; Loffler, Yourganov, Wilkinson, & Wilson, 2005).

In the face space, a particular identity trajectory originates at the average face (identity strength = 0) and terminates at the veridical face (identity strength = 1). Anti-caricatures (identity strength < 1), which lie between the average and the veridical face, are less distinctive versions of the veridical face. Anti-faces (identity strength < 0) lie on the other side of the average face and have opposite feature values by comparison to their veridical versions (Blanz, O'Toole, Vetter, & Wild, 2000; Leopold et al., 2001).

Four male faces in the computationally defined face space served as original faces. Anti-caricatures and anti-faces were generated by morphing the veridical face towards, and beyond, the average male face, respectively (see Fig. 1 for example stimuli). Anti-faces were created with—0.75 identity strengths, avoiding morphing artifacts at more extreme strengths. Anti-caricatures were generated at 0.10 and 0.35 identity strength levels. Rotated faces were created by rotating faces 30° to the right.

All face images were presented in color on a 20-in. monitor with a resolution of 1920 × 1200 pixels. Face images were presented in the center of the screen without a fixation point. Each face image subtended approximately 10° of visual angle horizontally and was viewed from a distance of 50 cm.

2.2. Design

Adaptation condition was varied as a within-participants factor. Three adaptation conditions were included (Fig. 2). Participants identified briefly presented frontal view anti-caricatures in three conditions: (a) without adaptation (*no-adaptation*), (b) following adaptation to a frontal anti-face (*within-view adaptation*), and (c) following adaptation to a rotated anti-face (*across-view adaptation*).

Familiarity was varied also, but as a between-participants factor. To the best of our knowledge, familiarity has not been manipulated previously in an adaptation study. Given the lack of data, the familiarity

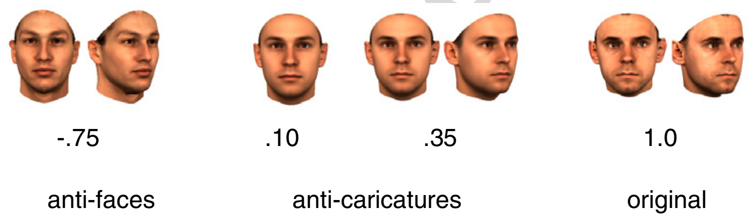

Fig. 1. Example stimuli illustrating the whole range of identity strengths used for a particular face (Alex) in the experiment.

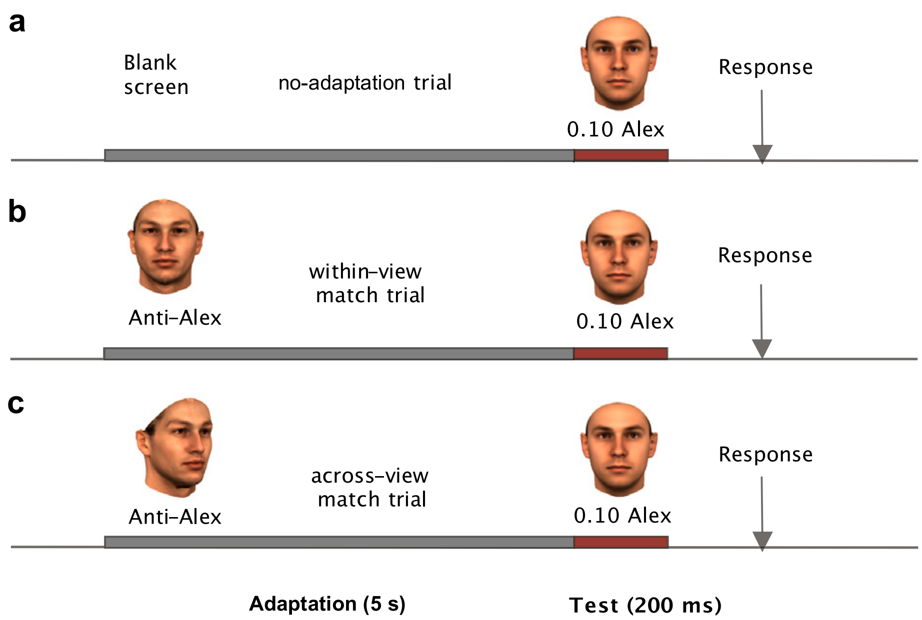

Fig. 2. Example trials. (a) No-adaptation trial. Participants identified briefly presented (200 ms) frontal 0.10 anti-caricatures without adaptation. (b and c) Adaptation trials in which identification task was performed following a 5-s period of adaptation to a frontal (within-view) or a rotated (across-view) anti-face.

Table 1
Familiarity condition

Familiarity level	Condition	Repetition time and presentation view
Low familiarity	LF	2 exposures of frontal original faces 2 exposures of frontal 0.35 anti-caricatures
Medium familiarity	MF	4 exposures of frontal original faces 4 exposures of frontal 0.35 anti-caricatures
High familiarity	HF	
Frontal view	HFFV	8 exposures of frontal original faces 8 exposures of frontal 0.35 anti-caricatures
Rotated view	HFRV	8 exposures of rotated original faces 8 exposures of rotated 0.35 anti-caricatures
Multiple views	HFMV	4 exposures frontal + 4 exposures rotated original faces 4 exposures frontal + 4 exposures rotated 0.35 anti-caricatures
Extreme familiarity	EF	
Multiple views	EFMV	16 exposures frontal + 16 exposures rotated original faces 16 exposures frontal + 16 exposures rotated 0.35 anti-caricatures

conditions we tested in current study were designed to be representative rather than exhaustive. We included six familiarity conditions varying the number of exposures and views of the familiarization stimuli (Table 1). In the *low familiarity* (LF) condition, participants saw each original frontal-view familiarization face twice. In the *medium familiarity* (MF) condition, participants saw the original frontal-view familiarization faces four times. This familiarity level is equivalent to the amount of familiarity participants had with faces in the Jiang et al. (2006) study. In the *high familiarity* (HF) condition, the total number of exposures to the original faces was set to eight. The HF condition was further divided into three familiarization conditions in which participants saw only frontal views (HFFV), only rotated views (HFRV), or half frontal and half rotated views (HFMV). Finally, in the *extreme familiarity multiple views* (EFMV) condition, participants saw the original faces 16 times from frontal view and 16 times from rotated view (see familiarization section for exact procedural details).

The effects of adaptation and familiarity were analyzed with an omnibus analysis of variance (ANOVA). Further planned comparisons were used to test for differences among the three high familiarity conditions. Specifically, we compared high familiarity with the rotated view with high familiarity with the frontal view, using a contrast between the HFRV and the HFFV conditions. We also tested the advantage of high familiarity with multiple views versus high familiarity with the frontal or rotated view alone, using a contrast between the HFMV and the combined HFFV/HFRV conditions.

2.3. Participants

A total of 102 undergraduate students from The University of Texas at Dallas participated in the study. The data from 12 (0.118) participants were discarded based on poor performance across all conditions in the test trials. This left 15 participants in each of the six familiarity conditions. All participants were naïve to the purpose of the study. Written consents were obtained before the experiment.

2.4. Procedure

The experiment started with a familiarization session, followed by a practice session and a test session.

2.4.1. Familiarization

The familiarization proceeded as follows. Participants were asked first to learn names for the four original faces. To do so, they viewed these four faces presented with their names, for 5 s each in random order. The exposures and views of the original faces varied according to the familiarity condition (see Table 1). For example, in the LF condition, a participant saw the four original faces, two times for 5 s each in random order. Then, with feedback given, participants were asked to name each original face once, by pressing one of the four appropriately labeled keys on the keyboard. Nearly all participants named the faces with no errors at this point.

Next, because the test faces in the adaptation study were 0.10 anti-caricatures (see below), we also familiarized participants with the 0.35 anti-caricature versions of the faces. These anti-caricatures have identity strengths between the original (1.0) and test face strengths (0.10), but are still easy to identify. The familiarization procedure was repeated, exactly, but with 0.35 anti-caricatures of the four original faces. Thus, a participant in the LF condition would see the 0.35 anti-caricatures of the four faces two more times, for 5 s each in random order. Finally, with feedback given, participants were asked to name each 0.35 anti-caricature once, by pressing one of the four appropriately labeled keys on the keyboard. There were virtually no errors at this point, and all participants claimed to know the names of the faces. No participants were eliminated due to poor performance in the familiarization session.

2.4.2. Practice

Practice and test sessions were identical for all participants, regardless of familiarity condition. Participants were given 40 practice trials in the format of the test trials, but with 0.35 anti-caricatures as test faces. These trials were meant to acquaint participants with the task and were excluded from the analysis.

2.4.3. Test

The test session began immediately after the practice session. Because of the large number of familiarity conditions, measuring adaptation as a function of identity strength (Jiang et al., 2006; Leopold et al., 2001) in each familiarity condition was impractical. Instead, we compared adaptation across familiarity conditions at a single, low identity strength (0.10). This identity strength has been shown to elicit consistent identity aftereffects in previous studies (Jiang et al., 2006; Leopold et al., 2001). This simplified the experimental design.

During test, participants identified briefly (200 ms) flashed, frontal views of 0.10 level anti-caricatures as one of the four original faces. In the no-adaptation trials, participants identified anti-caricatures without adaptation. In the within-view (across-view) adaptation trials, the identification task was performed following 5 s of adaptation to a frontal (rotated) anti-face. In "match" trials, the adapting and test stimuli were from same identity trajectories. In "non-match" trials, the adapting and test stimuli were from different identity trajectories. Non-match trials were added to prevent participants from learning pairs of anti-faces and anti-caricatures. Only the data from the match trials were included in the analysis.

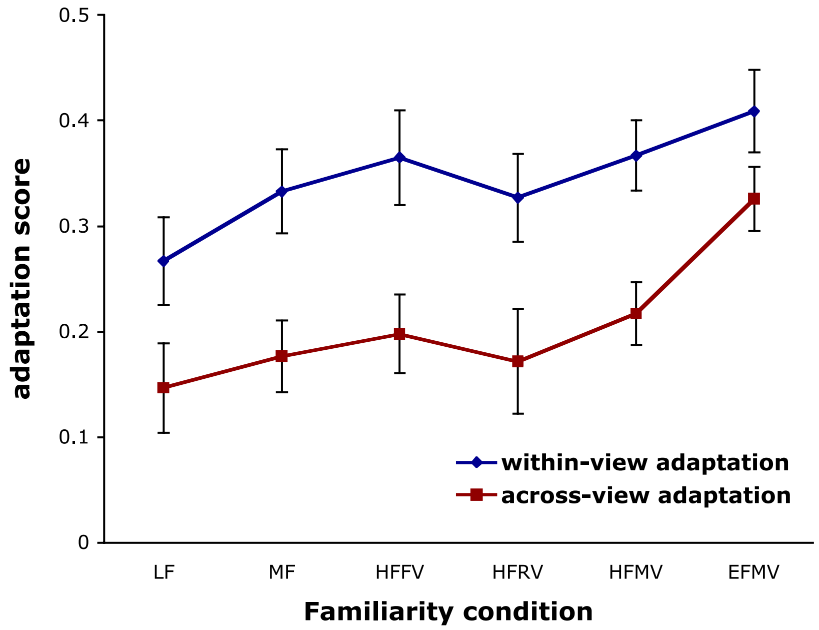

Fig. 3. Effects of familiarity on identity adaptation scores. Adaptation scores are computed as the difference between the identification accuracy in each adaptation condition and the no-adaptation control condition.

Each participant performed a total of 144 trials. Equal numbers of no-adaptation, within-view adaptation, and across-view adaptation trials were presented. For the within-view and across-view adaptation trials, equal numbers of match and non-match trials were included. Trial order was randomized for each participant, with the exception that every four no-adaptation trials were grouped together as one unit. This was to minimize the possible impact of a remaining aftereffect from previous adaptation trials.

3. Results

The proportion of correct identification of the test faces (i.e., the 0.10 anti-caricatures) was calculated for no-adaptation, within-view adaptation, and across-view adaptation conditions for each participant. These proportions were converted to "adaptation scores", which were computed for each individual participant as the difference between identification accuracy in each adaptation condition and the no-adaptation baseline condition. We used adaptation scores rather than proportion correct to control for possible increases in participant sensitivity to the familiarized faces (e.g., enhanced face naming skill) that might result from the familiarization process. These effects should influence performance in the no-adaptation condition as well and so this no-adaptation condition can be used to estimate, and parcel out, sensitivity effects in the adaptation conditions.[1]

As expected, identity adaptation effects occurred both within and across viewpoint, as indicated by adaptation scores that were significantly greater than zero throughout (see Fig. 3). Consistent with Jiang et al. (2006) and Jeffery et al. (2006), adaptation effects were stronger when the adapting and test faces were presented from the same view than when they were from different views, ($F(1, 84) = 106.87, p < .0001$). The magnitude of adaptation effects varied as a function of familiarity condition ($F(5, 84) = 2.33, p < .05$), with stronger adaptation effects when the faces were more familiar (see Fig. 3). It is worth noting that the adaptation effect for the across-view condition at the highest level of familiarity was statistically equivalent to the adaptation effect within-view for the least familiar faces (i.e., LF). This indicates an example of adaptation transfer for familiar faces that is equal to within-view adaptation for less familiar faces. No interaction was found between adaptation condition and familiarity level ($F < 1$).

Planned comparisons were used to test for differences among three high familiarity conditions, which included different views in the familiarization protocol. No difference was found between high familiarity with the rotated view and high familiarity with the frontal view, ($t(84) = .64$, ns). There was no significant advantage for familiarity with multiple views versus familiarity with the frontal or rotated views alone, ($t(84) = .59$, ns).

4. Discussion

The novel findings of this study are as follows. First, the overall magnitude of identity aftereffects increased as familiarity with faces increased. Second, familiarity with faces enhanced the transferability of identity adaptation effects over three-dimensional viewpoint change. Third, we found no evidence for an advantage for high familiarity with multi-

[1] Note that correct identification in the no-adaptation condition increased with familiarity, ($F(5, 84) = 3.98, p < .003$), which likely indicates sensitivity or naming facilitation with familiarity, supporting the use of adaptation score, rather than proportion correct, as an appropriate measure.

ple views over high familiarity with either the frontal or the rotated view alone. We discuss each of these findings, in turn.

4.1. Familiarity influences the magnitude of identity aftereffects

The present results reveal an important role for familiarity in strengthening identity aftereffects. This increase in the magnitude of identity aftereffects suggests that experience with a face may contribute to the development of a face representation that is more malleable and perceptually accessible than a newly formed face representation.

These adaptation results are consistent with previous findings that familiar faces exhibit greater perceptual tolerance to caricature manipulations (e.g., Benson & Perrett, 1991). Participants require some familiarity with faces to show caricature effects (e.g., Deffenbacher, Johanson, Vetter, & O'Toole, 2000) and tend to accept higher degrees of caricature distortion for familiar versus unfamiliar faces. Benson and Perrett (1991), for example, found that when choosing a good likeness of a person, participants selected more extreme caricatures for more familiar faces. In the present study, the difference in the magnitude of adaptation for familiar and less familiar faces indicates that the perceptual bounds of identity can be stretched more easily for familiar faces.

4.2. Familiarity influences three-dimensional transfer of identity aftereffects

We found that familiarity with faces enhanced the transferability of identity adaptation effects over three-dimensional viewpoint change. These results provide a novel type of "adaptation-based" evidence for the evolution of more robust face representations through experience. As a tool, adaptation has the potential to address fundamental questions about how we build view-invariant representations of faces and objects from view-constrained input.

The nature of face and object representations as view-specific or view-invariant is a long-standing issue, still in debate. The present findings suggest a re-focusing of this debate from a dichotomy between view-specific and view-invariant codes to a progression from view-specific to more robust view generalizable codes. At the highest level of familiarity we tested, the degree of adaptation transfer approached the magnitude of the within-view adaptation effects, nearly eliminating the advantage of adapting and testing from the same viewpoint. The mediating effects of familiarity on adaptation transfer suggest a caution in interpreting the size of adaptation transfer effects in absolute terms as evidence for particular kinds of representations (e.g., view-specific or view-invariant).

4.3. No evidence for a multiple-view advantage

Although we expected to find a significant advantage for high familiarity with multiple views over high familiarity with either the frontal or the rotated view alone, evidence for this was lacking. One possibility is that the inclusion of only two views (i.e., the frontal and the rotated views) in our multiple view condition was not sufficient to support an advantage. This issue can be addressed in future studies that systematically manipulate the number and range of familiarization views.

A second possibility is that multiple exposures to single views of faces can be at least partially effective in building a more view-transferable face representation. There are psychological and computational lines of evidence for this second possibility. On the psychological side, a recent study by Roark, O'Toole, Abdi, and Barrett (2006) demonstrates that repeated exposure to a single frontal image of a face increases accuracy for recognizing the person from a profile view or from a low resolution "gait video" (i.e., a person walking by a camera). This suggests that the generalizability of face representations develops even in the absence of additional information from exposure to new viewpoints.

On the computational side, Blanz and Vetter (1999) showed that it is possible to estimate novel views of an individual face by combining a two-dimensional image of the face with statistical data from a large number of three-dimensional laser-scanned heads. Using this approach, Blanz, Grother, Phillips, and Vetter (2005) synthesized frontal views of faces from 45° rotated face views. These synthesized views were given as input to "view-specific" face recognition algorithms (i.e., algorithms optimized for frontal view face recognition). Blanz et al. (2005) found that, for nine of the 10 face algorithms tested in the Face Recognition Vendor Test 2002 (Phillips, Micheals, Blackburn, Tabassi, & Bone, 2003), the use of synthesized frontal faces improved the rate of correct verification and identification over the use of the original profile views.

Combined, both the psychological and computational data suggest that experience with faces, in general, can support some view transfer for particular faces. The ability to use a lifetime of general knowledge about faces for building better representations of novel faces from view-constrained experience, does not exclude the possibility that exposure to additional views can enrich the robustness of face representations in a complementary way. This may occur via processes similar to those posited in view-based recognition theories (Poggio & Edelman, 1990; Riesenhuber & Poggio, 1999, 2000), where faces are represented by multiple view-specific mechanisms that become associated, over time, with experience. By this account, the degree of adaptation transfer across three-dimensional viewpoint change might reflect a strengthening of the connections among view-specific templates.

The present data suggest that a complete model of face representation might include a role for two types of learning processes, with one process exploiting the statistics acquired from general experience and a second process incorporating new views of individual faces into the representation. Combined, these two processes act to enhance view generalization for more familiar faces.

In summary, adaptation can serve as a useful tool for tracking the view generalizability of real face representations as they evolve. Using adaptation along with systematic manipulations of familiarity, it may be possible to address fundamental questions underlying the evolution of perceptual constancy for faces.

Acknowledgments

Thanks are due to Dana Roark and two anonymous reviewers for helpful comments on a previous version of this manuscript. We thank anonymous reviewer 2 for the suggestion of adaptation score as a measure of the familiarity effects.

References

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

Benson, P. J., & Perrett, D. I. (1991). Perception and recognition of photographic quality facial caricatures: implications for the recognition of natural images. *European Journal of Cognitive Psychology, 3*(1), 105–135.

Blanz, V., Grother, P., Phillips, J., & Vetter, T. (2005). Face recognition based on frontal views generated from non-frontal images. In *IEEE Conference on Computer Vision and Pattern Recognition CVPR*.

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*(8), 885–891.

Blanz, V. & Vetter, T. (1999). A morphable model for the synthesis of three-dimensional faces. In *Computer Graphics Proc. SIGGRAPH' 99*, pp. 187–194.

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.

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.

Deffenbacher, K. A., Johanson, J., Vetter, T., & O'Toole, A. J. (2000). The face typicality-recognizability relationship: encoding or retrieval locus? *Memory & Cognition, 28*, 1173–1182.

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*(4), 1128–1139.

Giese, M. A., & Leopold, D. A. (2005). Physiologically inspired neural model for the encoding of face spaces. *Neurocomputing, 65–66*, 93–101.

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

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

Jeffery, L., Rhodes, G., & Busey, T. (2006). View-specific coding of face shape. *Psychological Science, 17*(6), 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*(6), 493–500.

Kanwisher, N., McDermott, J., & Chun, M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for the perception of faces. *Journal of Neuroscience, 17*, 4302–4311.

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.

Loffler, G., Yourganov, G., Wilkinson, F., & Wilson, H. R. (2005). FMRI evidence for the neural representation of faces. *Nature Neuroscience, 8*(10), 1386–1390.

Logothetis, N., Pauls, J., Bülthoff, H., & Poggio, T. (1994). View-dependent object recognition by monkeys. *Current Biology, 4*, 401–411.

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

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

Phillips, P.J. Grother, P., Micheals, R., Blackburn, D., Tabassi, E., & Bone, J. M. (2003). FRVT 2002 evaluation report, Tech. Rep. NISTIR 6965. Available from http://www.frvt.org.

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

Pourtois, G., Schwartz, S., Seghier, M. L., Lazeyras, F., & Vuilleumier, P. (2005). View-independent coding of face identity in frontal and temporal cortices is modulated by familiarity: An event-related fMRI study. *Neuroimage, 24*, 1214–1224.

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.

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

Riesenhuber, M., & Poggio, T. (2000). Models of object recognition. *Nature Neuroscience, 3*(Supplement), 1199–1204.

Roark, D. A., O'Toole, A. J., Abdi, H., & Barrett, S. (2006). Learning the moves: the effect of familiarity and facial motion on person recognition across large changes in viewing format. *Perception, 35*, 761–773.

Suzuki, S., & Cavanagh, P. (1998). A shape-contrast effect for briefly presented stimuli. *Journal of Experimental Psychology: Human Perception and Performance, 24*(5), 1–27.

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

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

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 after-effects in the perception of faces. *Psychonomic Bulletin Review, 6*, 647–653.

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