

Figural aftereffects in the perception of faces

MICHAEL A. WEBSTER and OTTO H. MACLIN
University of Nevada, Reno, Nevada

We examined figural aftereffects in images of human faces, for which changes in configuration are highly discriminable. Observers either matched or rated faces before or after viewing distorted images of faces. Prior adaptation strongly biases face perception by causing the original face to appear distorted in a direction opposite to the adapting distortion. Aftereffects transferred across different faces and were similar for upright or inverted faces, but were weaker when the adapting and test faces had different orientations (e.g., adapt inverted and test upright). Thus the aftereffects depend on which images are distorted, and not simply on the type of distortion introduced. We further show that the aftereffects are asymmetric, for adapting to the original face has little effect on the perception of a distorted face. This asymmetry suggests that adaptation may play an important normalizing role in face perception. Our results suggest that in normal viewing, figural aftereffects may strongly influence form perception and could provide a novel method for probing properties of human face perception.

Visual perception can be strongly affected by changes in the observer's state of adaptation. Looking at patterns induces a variety of related aftereffects, including changes in detection thresholds, in the apparent contrast of suprathreshold stimuli, and in the appearance of stimuli (Webster, 1996). In this study, we focus on the changes in perceived shape, or figural aftereffects, that result from adaptation to natural images. Classical examples of figural aftereffects include changes in apparent size (Köhler & Wallach, 1944), orientation (Gibson & Radner, 1937), curvature (Gibson, 1933), and spatial frequency (Blakemore & Sutton, 1969). Analogous aftereffects occur for perceptual dimensions other than form (e.g., motion aftereffects, Wohlgenuth, 1911, or changes in perceived hue following color contrast adaptation, Webster & Mollon, 1994), suggesting that such adjustments are a general and important feature of sensory coding (Barlow, 1990). Aftereffects are induced with only brief exposures and can be readily induced with complex, natural images (Köhler & Wallach, 1944; as, e.g., in the "waterfall illusion," Addams, 1834). Yet the patterns of adaptation that might routinely arise in natural viewing, and how these adaptational adjustments might influence normal visual perception, have received little attention. Webster and Mollon (1997) and Webster and Miyahara (1997) found large changes in color and spatial sensitivity that result from adaptation to the color and spatial statistics of natural images, and suggested that these adjustments are an intrinsic part of the visual response to any image. Here we examine the implications of these natural sensitivity adjustments for suprathreshold form perception.

To examine figural aftereffects for natural images, we measured the aftereffects that are induced in the appearance of images of human faces. We chose faces as stimuli because, compared with other classes of natural images, faces appear highly discriminable and easy to identify, so that observers should be highly sensitive to any changes induced by adaptation. Further, the capacity for face recognition appears to depend strongly on sensitivity to higher-order configural properties of the stimulus (Diamond & Carey, 1986; Rhodes, 1988; Tanaka & Farah, 1993) rather than the perception of individual features, and may therefore reveal properties of form perception that are more relevant to the perception of complex images.

We also focused on faces as natural stimuli because they have obvious biological significance and may depend on distinct (if not face-specific) neural mechanisms. Several lines of evidence suggest that face recognition involves specialized processes and pathways in the brain. Behaviorally, face recognition can be strongly disrupted by inverting the orientation of the image (Valentine, 1988; Yin, 1969) or by inverting image contrast (Galper, 1970). Because these manipulations preserve all of the information in the image, such inversion effects demonstrate that recognition does not depend on a generic representation of spatial information. Mechanisms specific for face coding have also been suggested by demonstrations of "face-selective" responses in neurons in a variety of areas in the primate brain (Desimone, 1991; Logothetis & Sheinberg, 1996; Perrett, Mistlin, & Chitty, 1987; Rolls & Baylis, 1986). Finally, localized cortical processing specific for face coding is suggested both by the clinical deficits of prosopagnosia (Damasio, 1985; Farah, Wilson, Drain, & Tanaka, 1995) and by imaging studies (Kanwisher, McDermott, & Chun, 1997; Sergent, Ohta, & MacDonald, 1992). Collectively such results suggest that in human vision a specialized organization exists for face coding, though the nature of this code remains elusive. The prop-

The work reported here was supported by Grant EY10834 from the National Eye Institute. Correspondence should be addressed to M. A. Webster, Department of Psychology/296, University of Nevada, Reno, NV 89557 (e-mail: mwebster@scs.unr.edu).

erties of figural aftereffects for faces might reflect characteristics of this code.

METHOD

Stimuli

We measured how the appearance of a face is influenced by prior exposure to a distorted version of the original face. This allowed us to examine adapting and test stimuli that differed in known and quantified ways. Face images were taken from the neutral-expression face set of Matsumoto and Ekman (1988). The images were distorted by expanding or contracting the frontal-view image of the face relative to a midpoint on the nose. The distortion was weighted by a circular Gaussian envelope so that the largest changes were near the midpoint and so that there were only small changes in the outline of the head. Images were created by remapping pixel values from a location within the original image (x_j, y_j) to a given location in the new image (x_i, y_i) . For x values,

$$x_j = (1 - w)x_i + wx_c,$$

where x_c equals the center point

and

$$w = \alpha \exp\{-[(x_i - x_c)^2 + (y_i - y_c)^2]/2\sigma^2\},$$

where $\sigma = 0.2$ times the head width.

The sign of w was changed for pixels to the left or right of center. Similar equations describe the remapping of y values. By varying the amplitude (α) of the distortion (from -1 to $+1$ in steps of $.04$ for most experiments), we generated large (51×51) arrays of images whose two dimensions corresponded to varying degrees of contraction or expansion along the horizontal or vertical axes of the image. Figure 1 shows an example of the resulting stimulus array. The images show distortions in steps of 0.5 , with the original image at the center.

Procedure

During experiments only a single test image was presented on the display monitor at any time (for 1 sec, interleaved with a blank screen or an adapting image). We used two different tasks to measure the aftereffects.

Nulling Matches

In the first task, subjects learned to recognize the original, undistorted image and then chose the image from the array that appeared most like the original, before or after adaptation. Training for the original face was given at the beginning of each session. After inspecting the target face (for up to several seconds), a random image from the array was displayed. Subjects then tried to adjust the image back to the original by using two pairs of buttons to scroll through the array coordinates defining the image. (That is, each buttonpress increased or decreased the current x [for one button pair] or y [for the second button pair] index to the image array, which defined the next displayed image. Up to several buttonpresses could be made during the interval between each test image presentation.) Feedback was given for each match. This training sequence continued until observers were able to find the original (within a 0.1 distortion) 10 times when a single target face was tested or 5 times each when multiple target faces were tested (as in Figures 2 and 4). Next, in a preadapt testing phase, they were again presented a randomly chosen image and were instructed to vary the displayed image to try to find the closest match to the original face, but in this case with no further feedback or opportunities for reinspecting the original. Finally, in the adapt phase, this procedure was repeated after viewing (continuously but without specific fixation instructions) a single adapting image presented for 5 min, and then for 8-sec intervals interleaved with each presentation of the test image (with 0.5-sec uniform-gray screens between the test and adapt images). (Although we did not formally test the time course of the aftereffects, we observed infor-

mally that the adaptation built up and decayed fairly rapidly, as typical of figural aftereffects. The readaptation intervals were therefore included to maintain a constant state of adaptation. Preadapt testing followed an identical timing sequence but with uniform-gray screens in place of the adapting images.) During a daily session only a single adapting distortion was tested. Observers made 10 preadapt and 10 postadapt settings for sessions measuring a single test face, or 5 randomly interleaved settings for each test face in sessions with multiple target faces. Observers in this task included the authors and Subjects A.R. and V.R., who were unaware of the specific aims of the experiment.

Normal-Distort Ratings

We verified the pattern of adaptation effects revealed by the matching task by using an alternative procedure in which observers were asked to simply rate whether the displayed image appeared normal or distorted, rather than set a match to a particular face image. In this case observers received no training or information to indicate which image from the array was the original. During testing, the image displayed was varied by using eight randomly interleaved staircases that tracked the ratings along different axes of the array (at 45° intervals). That is, for each axis a "normal" response caused the next image for that axis to be more distorted, and a "distorted" response caused the next image to be less distorted. The mean of the normal ratings was estimated by averaging the eight staircase-reversal points (i.e., the points along each axis at which "normal" and "distorted" ratings were equally likely). Ratings were made before or after adapting to a single face image, following the same timing sequence as in the matching task. For the rating task, measurements were made for 27 naive observers for the upright face images. Each observer participated in only a single session and was tested with a single adapting distortion, with 3 different observers tested for each of nine different adapting images (the original image or one of eight distorted images). A further 24 observers made settings with inverted test or adapt images (again with 3 different observers tested for each specific adapting and testing condition). All figures plot the settings averaged across observers.

RESULTS

Aftereffects for Distorted Adapting Images

Figure 1 shows three examples of the measured aftereffects for face adaptation. Each shows that adaptation to a distorted face causes the original face to appear distorted in a direction *opposite* to the adapting distortion. The measurements reveal this aftereffect indirectly by showing that the stimuli required to null the aftereffect are physically distorted in a direction *toward* the adapting distortion. For example, after adapting to an expanded face, the original face appeared too contracted. Thus to null this perceptual contraction, observers had to choose a face that was physically expanded.

The left panel of Figure 1 illustrates this result by showing how the face matches were affected by adapting to the original face or eight different directions of adapting distortion (unfilled diamonds). In the preadapt phase (or following adaptation to the original image), the settings are close to the original image (unfilled circles). However, the nulls are systematically shifted toward each distorted adapting image following the adaptation to these images (filled triangles). Comparable perceptual shifts occurred in the rating task, as shown in the middle panel of Figure 1. Before adaptation, the mean of the images rated as normal is close to the original, undistorted image

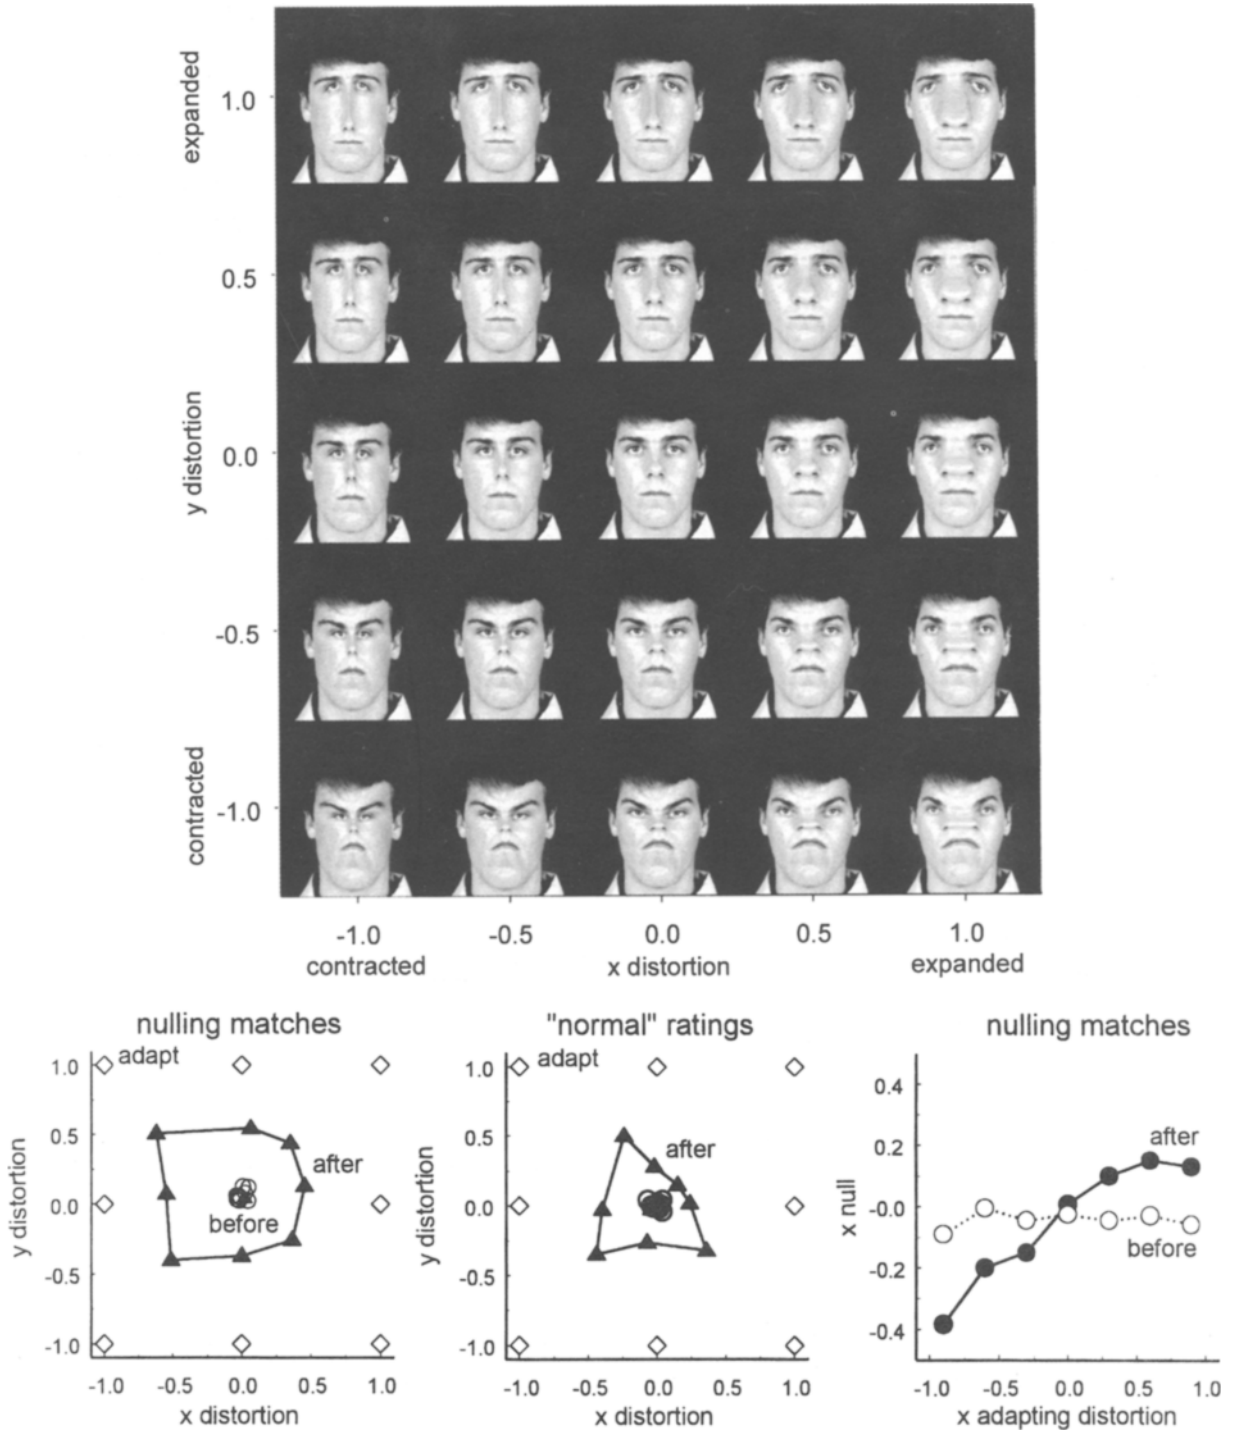


Figure 1. An example of the face array and aftereffects induced by viewing distorted images in the array, measured by nulling matches or “normal versus distorted” ratings. Points plotted show the settings before (unfilled circles) and after (filled symbols) adapting to distorted images (with coordinates shown by the unfilled diamonds). Adapting to any direction of distortion causes the original face to appear biased in the opposite direction, and must be nulled by choosing images shifted toward the adapting distortion.

(unfilled circles). Yet the normal ratings after adaptation (filled triangles) are shifted toward each adapting image, again indicating that the adaptation caused the original image to appear distorted in the opposite direction.

Finally, the right panel of Figure 1 illustrates how the aftereffects depend on the magnitude of the adapting distortion. The figure plots the nulling matches made before (unfilled symbols) or after (filled symbols) adapting

to different degrees of contraction or expansion along the horizontal (x) axis of the images. Again, adaptation produces systematic shifts in the settings along the adapting (x) axis, with little shift along the second (y) axis of the array. The results illustrate that aftereffects occur both for small distortions—that like the original face appear as plausible images of “real” faces under neutral adaptation—and for distortions so large that they do not appear as a possible face (though, as our rating results suggest, such judgments may depend strongly on the state of adaptation).

The aftereffects illustrated in Figure 1 are perceptually salient and reflect actual shifts in the perception—rather than simply the memory—of the test image. For example, whereas in the matching task, subjects were required to set images according to a remembered target face, in the rating task they were never informed about the original face, yet the pattern of aftereffects was similar. Further, we observed that different aftereffects could be induced simultaneously in the same images presented to two retinal regions (to either side of fixation) that were adapted to different images, confirming the retinotopic specificity characteristic of figural aftereffects. (We could not adopt this stimulus arrangement for measuring the aftereffects because recognition was sharply degraded for peripherally viewed images.)

Adaptation Effects for Different Faces

Aftereffects were similar for different face images, as confirmed for five different images from the Matsumoto and Ekman (1988) set (Face A shown in Figure 1 and Faces B–E in Figure 2), as well as for a number of other face images that varied from previously unfamiliar to highly familiar (e.g., the observer’s own face). The aftereffects also exhibited strong transfer across adapting and test stimuli drawn from different faces. For example, Figure 2 shows how adaptation to Face A distortions affected the matches for Faces B–E. The matches again indicate that each test face appeared biased in appearance away from the adapting distortion (filled circles), and in some cases these biases were as large as for the same-face adaptation (unfilled triangles). Notably, the measurements do not show a shift in the matches for other faces after adapting to the original face (filled triangles), a result that could be related to the asymmetries in adaptation between normal and distorted faces (see below). However, our stimuli could vary only along the dimensions of our distortions, and thus might not reveal aftereffects within the (unknown) dimensions along which two normal faces differ.

Effects of Image Orientation

An upside-down face is more difficult to recognize, and large distortions in inverted faces can often go unnoticed

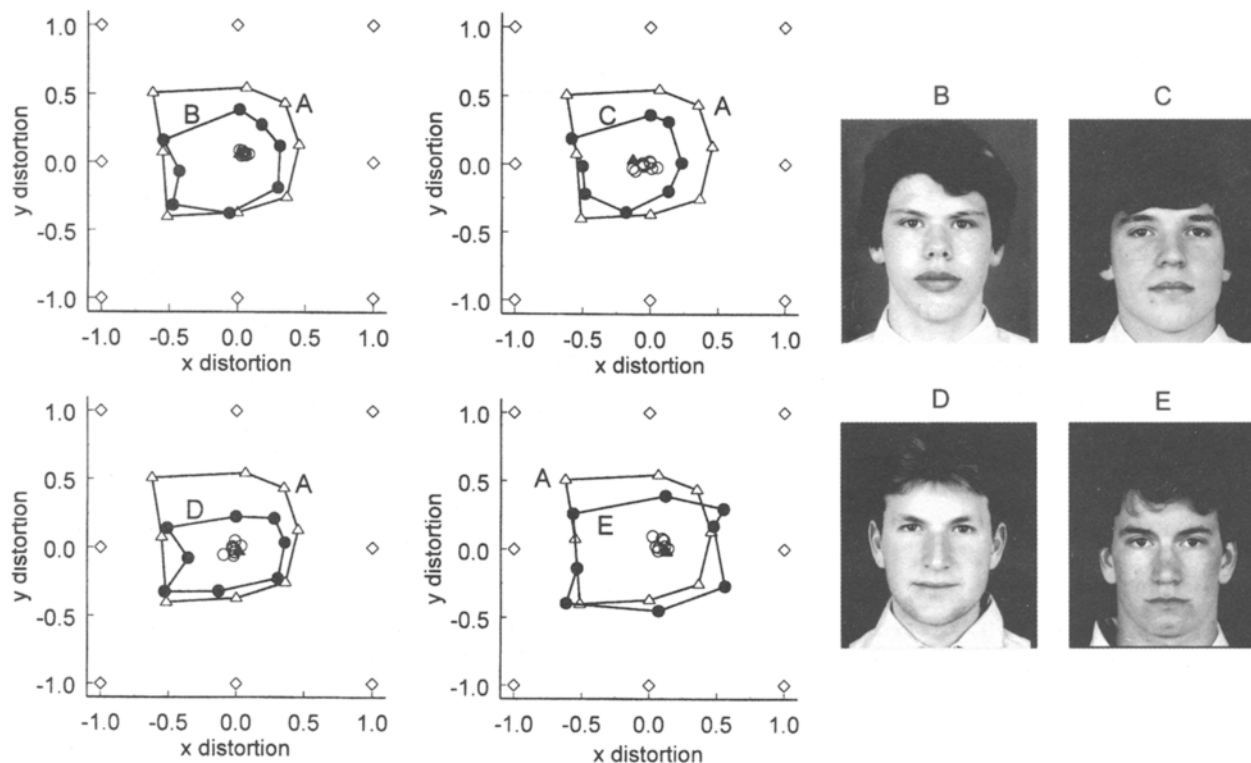


Figure 2. Effects of adapting to distortions in the Face A image (from Figure 1) on the nulling matches for four different face images (Faces B–E). Each panel shows matches for different-face adaptation (e.g., adapt to A, test B; filled circles), compared with the same-face adaptation (adapt to A, test A; unfilled triangles). Unfilled diamonds show the coordinates of the adapting distortions.

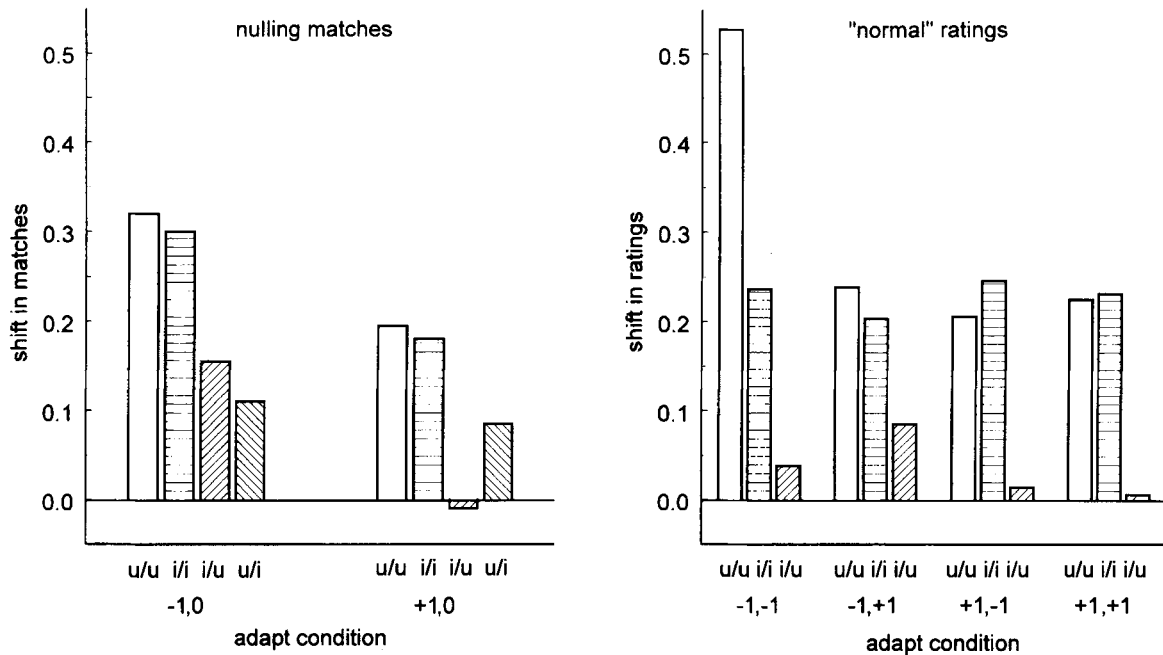


Figure 3. Effects of image orientation on the aftereffects. Histograms show the magnitude of the shift (toward the adapting distortion) for upright-adapt/upright test (u/u), inverted-adapt/inverted-test (i/i), inverted-adapt/upright-test (i/u), or upright-adapt/inverted-test (u/i). Left panel shows results from the matching task; right panel shows results from the rating task.

(see, e.g., Thompson, 1980). As noted, this inversion effect suggests that face recognition depends on specialized perceptual processes. We tested for an inversion effect in face adaptation by comparing the aftereffects when either the test face or adapt face was upright or inverted. In the matching task, the adapting images had horizontal distortions of -1 or $+1$. As shown in Figure 3, the mean shifts in the nulls were comparable whether both images were upright or both were inverted, suggesting that there is not a conventional inversion effect for the figural aftereffects (i.e., the poorer recognition associated with inverted images did not result in a weaker figural aftereffect). On the other hand, aftereffects were substantially weaker when the adapt and test images had *different* orientations (i.e., adapt inverted and test upright or adapt upright and test inverted). These differences were assessed with a three-way analysis of variance (ANOVA) (adapt distortion \times test orientation \times adapt orientation), which showed significant effects only for adapt distortion [$F(1,23) = 6.3, p < .05$] and for the interaction between test and adapt orientation [$F(1,23) = 16.0, p < .001$]. The weaker effect when the adapt and test have different orientations suggests that the adaptation is not merely to the distortion gradient in the image (which is identical for the upright and inverted faces), but depends instead on the specific spatial configuration of the images. These effects were again confirmed with the rating task (Figure 3). In this case, measurements were made for four adapting distortions ($+1, +1$; $+1, -1$; $-1, +1$; or $-1, -1$). The ratings again revealed comparable shifts when the test and adapt had the same orientation (averaging 0.29 for

both-upright vs. 0.23 for both-inverted), but were clearly weaker when the orientations differed (0.05 for adapt inverted and test upright).

Adaptation to Distorted Versus Undistorted Images

As the preceding results show, adaptation to distorted faces can strongly bias the appearance of undistorted faces. Yet, conversely, we found that adapting to the original face has little effect on the appearance of distorted images. This asymmetry is illustrated in Figure 4. Observers adapted to each of the five images shown in Figure 4 and then set nulling matches for each of the remaining four images. Adaptation to each distorted image produced large and comparable shifts away from the adapting image in the appearance of each test image (i.e., both the undistorted and the three other distorted test images), and thus large nulling shifts in the matches toward each adapting image. In contrast, adaptation to the undistorted image (N) induced little aftereffect (middle histogram of Figure 4). This asymmetry was confirmed by comparing how much the distorted adapting images shifted the matches for the undistorted test image ($M = 0.24$) or vice versa ($M = 0.04$; $t = 6.2, p < .001$).

The asymmetry between distorted and undistorted adapting images distinguishes face adaptation from the adaptation aftereffects typical of many stimuli, including size (Köhler & Wallach, 1944), spatial frequency (Blake-more & Sutton, 1969), auditory frequency (von Békésy, 1960), or hue (Webster & Mollon, 1994). For such stimuli, adaptation to any value causes stimuli of greater or lesser value to appear shifted away from the adapting value,

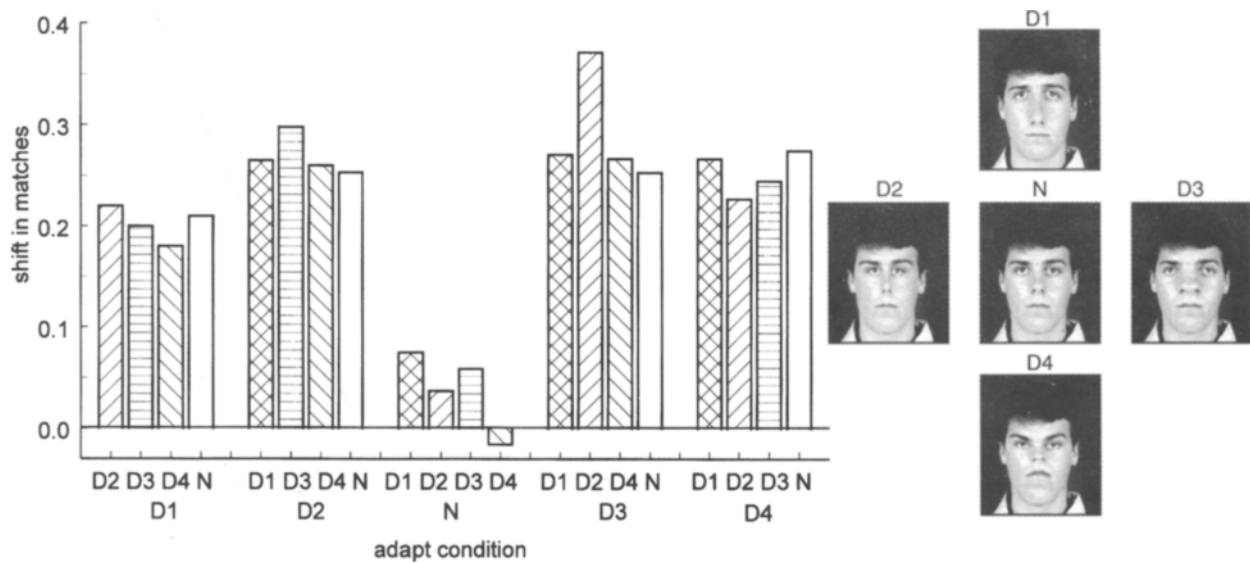


Figure 4. Asymmetries in the adaptation effects for the undistorted (N) image versus four distorted images (D1–D4, with x or y distortion amplitudes of ± 0.64 ; shown in the figure). The five sets of histograms plot the shift (toward each adapting distortion) that adaptation to one of the five images induced in the matches for the remaining four images. Adapting to any of the distorted images biases the matches of the original and the remaining distorted images, whereas adaptation to the original image (center histogram) has little effect on the perception of the distorted images.

while the adapting stimulus itself appears unchanged (the “distance paradox” of figural aftereffects; Köhler & Wallach, 1944; Osgood & Heyer, 1952; von Békésy, 1960). Such behavior implies a continuous dimension along which no value is special. That the figural aftereffects for faces do not follow this pattern rules out simple featural properties of the images (e.g., size or absolute eye separation) as the relevant stimuli for the adaptation.

Instead, the asymmetry suggests that perceptually, the images in the face array do not form a uniform continuum. Some faces (which may include or are similar to the original faces) have a special status and may represent a neutral point (or neutral range) for the adaptation, and the pattern of aftereffects appears to reflect a “renormalization” of this neutral point. Consistent with this, we observed that during adaptation the distorted adapting face itself gradually *appeared* less distorted over time. This shift was evident in the rating task, for some observers rated even a strongly distorted adapting image as normal after adaptation, though the same image was far outside the gamut of images rated as normal before adaptation.

This renormalization is characteristic of the figural aftereffects for patterns such as tilted (Gibson & Radner, 1937) or curved (Gibson, 1933) lines, which appear less tilted or less curved with prolonged viewing, and for which “vertical” or “straight” represents neutral percepts (though adaptation to orientation and curvature in line stimuli may exhibit little asymmetry in the aftereffects, Köhler & Wallach, 1944). The pattern of face aftereffects is also analogous to the changes in color appearance induced by chromatic (von Kries) adaptation. For example,

“white” has a neutral or special status in color space, and chromatic adaptation alters perceived colors by shifting the appearance of the adapting color toward this neutral point, with concomitant shifts in the appearance of all chromatic stimuli (Webster, 1996). The similarities in the form of adaptation for color and faces suggest that similar coding principles could underlie their representations.

DISCUSSION

To summarize, we have found that adaptation to a distorted face induces large changes in the appearance of an undistorted face. These aftereffects are induced by free viewing naturalistic stimuli for only brief periods of time. This suggests that in natural viewing, pattern adaptation may normally exert an important influence on form perception. Face images offer unique advantages for studying this influence because observers are so highly sensitive to their configural properties. In turn, such adaptation effects may be a prevalent factor affecting the coding and recognition of faces, and could thus be important for understanding some of the properties and dynamics of face coding.

While strongly affecting face perception, the adaptation effects themselves need not reflect processes specific to face recognition. However, the asymmetries in the aftereffects between normal and distorted faces suggest that the properties of adaptation for faces may differ from the characteristics of adaptation for other classes of stimuli (e.g., simple feature dimensions) and could thus reflect processes that are central to (if not specialized for) face perception. By measuring aftereffects for images

that capture more accurately the natural patterns of variation across faces, adaptation could provide a potential tool for exploring the perceptual dimensions underlying human face recognition (see, e.g., Atick, Griffin, & Redlich, 1996).

REFERENCES

- ADDAMS, J. (1834). An account of a peculiar optical phaenomenon seen after having looked at a moving body, etc. *London and Edinburgh Philosophical Magazine & Journal of Science*, **5**, 373-374.
- ATICK, J. J., GRIFFIN, P. A., & REDLICH, A. N. (1996). The vocabulary of shape: Principal shapes for probing perception and neural response. *Network: Computation in Neural Systems*, **7**, 1-5.
- BARLOW, H. B. (1990). A theory about the functional role and synaptic mechanism of visual after-effects. In C. Blakemore (Ed.), *Vision: Coding and efficiency* (pp. 363-375). Cambridge: Cambridge University Press.
- BLAKEMORE, C., & SUTTON, P. (1969). Size adaptation: A new aftereffect. *Science*, **166**, 245-247.
- DAMASIO, A. R. (1985). Prosopagnosia. *Trends in Neurosciences*, **8**, 132-135.
- DESIMONE, R. (1991). Face-selective cells in the temporal cortex of monkeys. *Journal of Cognitive Neuroscience*, **3**, 1-8.
- DIAMOND, R., & CAREY, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, **115**, 107-117.
- FARAH, M. J., WILSON, K. D., DRAIN, H. M., & TANAKA, J. R. (1995). The inverted face inversion effect in prosopagnosia: Evidence for mandatory, face-specific perceptual mechanisms. *Vision Research*, **35**, 2089-2093.
- GALPER, R. E. (1970). Recognition of faces in photographic negative. *Psychonomic Science*, **19**, 207-208.
- GIBSON, J. J. (1933). Adaptation, after-effect and contrast in the perception of curved lines. *Journal of Experimental Psychology*, **16**, 1-31.
- GIBSON, J. J., & RADNER, M. (1937). Adaptation, after-effect and contrast in the perception of tilted lines. I. Quantitative studies. *Journal of Experimental Psychology*, **20**, 453-467.
- 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.
- KÖHLER, W., & WALLACH, H. (1944). Figural aftereffects: An investigation of visual processes. *Proceedings of the American Philosophical Society*, **88**, 269-357.
- LOGOTHETIS, N. K., & SHEINBERG, D. L. (1996). Visual object recognition. *Annual Review of Neuroscience*, **19**, 577-621.
- MATSUMOTO, D., & EKMAN, P. (1988). *Japanese and Caucasian facial expressions of emotion (JACFEE) and neutral faces (JACNeuf)*. San Francisco: San Francisco State University, Department of Psychology.
- OSGOOD, C. E., & HEYER, A. W. (1952). A new interpretation of figural after-effects. *Psychological Review*, **59**, 98-118.
- PERRETT, D. I., MISTLIN, A. J., & CHITTY, A. J. (1987). Visual neurones responsive to faces. *Trends in Neurosciences*, **10**, 358-364.
- RHODES, G. (1988). Looking at faces: First-order and second-order features as determinants of facial appearance. *Perception*, **17**, 43-63.
- ROLLS, E. T., & BAYLIS, G. C. (1986). Size and contrast have only small effects on the responses to faces of neurons in the cortex of the superior temporal sulcus of the monkey. *Experimental Brain Research*, **65**, 38-48.
- SERGENT, J., OHTA, S., & MACDONALD, B. (1992). Functional anatomy of face and object processing. A positron emission topography study. *Brain*, **115**, 15.
- TANAKA, J. R., & FARAH, M. J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology*, **46A**, 225-245.
- THOMPSON, P. (1980). Margaret Thatcher: A new illusion. *Perception*, **9**, 483-484.
- VALENTINE, T. (1988). Upside-down faces: A review of the effect of inversion upon face recognition. *British Journal of Psychology*, **79**, 471-491.
- VON BÉKÉSY, G. (1960). *Experiments in hearing*. New York: McGraw-Hill.
- WEBSTER, M. A. (1996). Human colour perception and its adaptation. *Network: Computation in Neural Systems*, **7**, 587-634.
- WEBSTER, M. A., & MIYAHARA, E. (1997). Contrast adaptation and the spatial structure of natural images. *Journal of the Optical Society of America A*, **14**, 2355-2366.
- WEBSTER, M. A., & MOLLON, J. D. (1994). The influence of contrast adaptation on color appearance. *Vision Research*, **34**, 1993-2020.
- WEBSTER, M. A., & MOLLON, J. D. (1997). Adaptation and the color statistics of natural images. *Vision Research*, **37**, 3283-3298.
- WOHLGEMUTH, A. (1911). On the after-effect of seen movement. *British Journal of Psychology Monograph Supplements*, No. 1, 1-117.
- YIN, K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, **81**, 141-145.

(Manuscript received August 5, 1998;
revision accepted for publication March 17, 1999.)