Influence of adaptation on the perception of distortions in natural images

Otto H. MacLin
Michael A. Webster
University of Nevada, Reno
Department of Psychology
Reno, Nevada 89557
E-mail: mwebster@unr.nevada.edu

Abstract. *We examined judgments of distortions in images of faces and how they are influenced by visual adaptation to distorted images. Two-dimensional image arrays were generated by varying the degree of local expansion or contraction along the horizontal or vertical axis of the image. Images along different directions within the array were varied in staircases while observers rated the images as either "normal" or "distorted." The reversal points for the staircases define the gamut of distortions that observers accept as a normal variant of the original image, and they are well described by ellipses. Under neutral adaptation this gamut is roughly centered on the original image. However, after viewing a distorted face, the images perceived as normal are strongly biased toward the adapting distortion, and the range of the gamut increases along the axis of the distortion. Thus changes in the state of adaptation can markedly alter the perception of images, by altering both the perceived neutral point and the sensitivity to variations in images.* © 2001 SPIE and IS&T. [DOI: 10.1117/1.1330573]

1 Introduction

Many classic visual after-effects provide striking testament to the malleability of perception. For example, in the motion after-effect, viewing a drifting pattern causes a static image to appear to be moving in the opposite direction.[1] Analogous examples in form perception include figural after-effects such as the tilt after-effect (in which adaptation to oblique lines causes an upright line to appear tilted away from vertical in the opposite direction of the adapting stimulus[2]) and size after-effects (in which adaptation to patterns of a particular size or periodicity alters the apparent size[3] or spatial frequency[4] of subsequently viewed patterns), and the after-effects induced by sensory-motor adaptation to optical distortions.[5] The salience of such after-effects illustrates that visual coding can be strongly biased by changes in the state of adaptation. However, these perceptual changes have often been treated as illusions or anomalies of perception—induced under conditions that the visual system was not designed to operate in—rather than as a reflection of processes that are intrinsic to and necessary for perception (see, e.g., Ref. 6). Yet this interpretation is countered by the sheer prevalence of visual after-effects, and by the fact that they can be readily induced in natural viewing (as, e.g., in the waterfall illusion[7]). These observations suggest instead that the visual system is always operating under the influence of pattern adaptation, just as it is always operating in a state of light adaptation, and that these adaptation states are important in defining the "normal" properties of perception.[8]

We have previously examined[9] the figural after-effects that arise in the course of normal viewing of natural images by focusing on images of human faces and how their appearance is altered by adaptation. We chose faces as stimuli because human observers are extremely sensitive to facial configurations, as evidenced by a remarkable capacity to recognize an individual face and to discriminate between individuals, even though compared to other classes of natural images, human faces are physically highly similar. Observers should thus be sensitive to any changes in perceived configuration induced by the adaptation. We also examined faces because they are biologically important images, and may be encoded by specialized processes in the visual system (see, e.g., Ref. 10). Measurements of adaptation effects for faces may reveal important aspects of these processes, and may also be important for understanding the dynamics of face recognition in real-world contexts (e.g., in the many situations in which different faces are viewed in succession, and thus in which the appearance of an individual face may depend on the characteristics of faces viewed previously).

To explore figural after-effects for faces, we distorted images of a face in different ways by locally expanding or contracting the image along different axes of the image. We then asked how prior viewing of those distortions influenced the appearance of the face. This procedure had the advantage that the adapting and test stimuli were drawn from a single image and thus differed in known and quantified ways, and revealed surprisingly large influences on face perception. Specifically, viewing a distorted image of a face caused the original face to appear strongly distorted in the opposite direction.[9] Thus, for example, adaptation to a face image that was expanded caused the original to appear too contracted, and vice versa. These after-effects show that face perception and recognition can depend strongly on the observer's state of adaptation.

Paper HVEI-06 received July 21, 1999; revised manuscript received June 26, 2000; accepted Sep. 9, 2000.
1017-9909/2001/$15.00

Our earlier study[9] examined how adaptation influenced recognition of a single test image by examining which physical distortion was required to null the after-effect (e.g., by choosing the test image that looked most like the learned original image, or which image appeared most undistorted). In the present study we instead directly characterized how adaptation to images influences judgments of image distortions by measuring the *collection* of images that appeared normal or distorted before or after adaptation. These measurements allowed us to define the gamut of images (i.e., the range of distortions) that an observer would accept as normal variations in a stimulus. We use this procedure to measure acceptable variations in face images and to compare these to images from natural outdoor scenes. The properties of this gamut and how they change with adaptation have important implications for understanding the consequences of visual adaptation, and are relevant for understanding how adaptation may bias an observer's judgments of image quality. (A preliminary account of this work was given in Ref. 11.)

2 Methods

Stimuli were gray-level images of human faces or natural outdoor scenes. The frontal-view face images were scanned from slides taken from the neutral-expression face set of Matsumoto and Ekman.[12] Outdoor images were taken from the scenes used in the study of Webster and Miyahara.[13] Each image was spatially distorted by expanding or contracting the image relative to a midpoint (e.g., on the nose for a face). The magnitude of the distortion was weighted by a Gaussian envelope so that changes were largest near the midpoint and then tapered so that there was little change in the outline of the head. The same distortions were applied to the outdoor scenes. Note that these distortions are simple and arbitrary, and differ, for example, from distortions specifically designed to modulate properties of faces, as in caricaturing algorithms (see, e.g., Refs. 14 and 15). Nevertheless, the distortions we used alter the perceived configurations of the faces in ways that appear natural. For example, the horizontal distortions primarily appear to alter eye separation and nose width, while distortions along the vertical dimension primarily appear to alter the length of the nose and the eye-mouth separation (as seen in the horizontal and vertical axes of Fig. 1). And again, because we tapered the distortions, for modest values they appear more as configural changes in facial features than as a stretching or compression of the image per se.

The images were created by remapping pixel values from a location in the original image $(x_{\text{orig}}, y_{\text{orig}})$ to a location in the new image (x_i, y_i), according to

$$x_{\text{orig}} = (1 - w_x)x_i + w_x x_c$$

and

$$y_{\text{orig}} = (1 - w_y)y_i + w_y y_c,$$

where x_c, y_c equals the midpoint for the distortion,

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

and

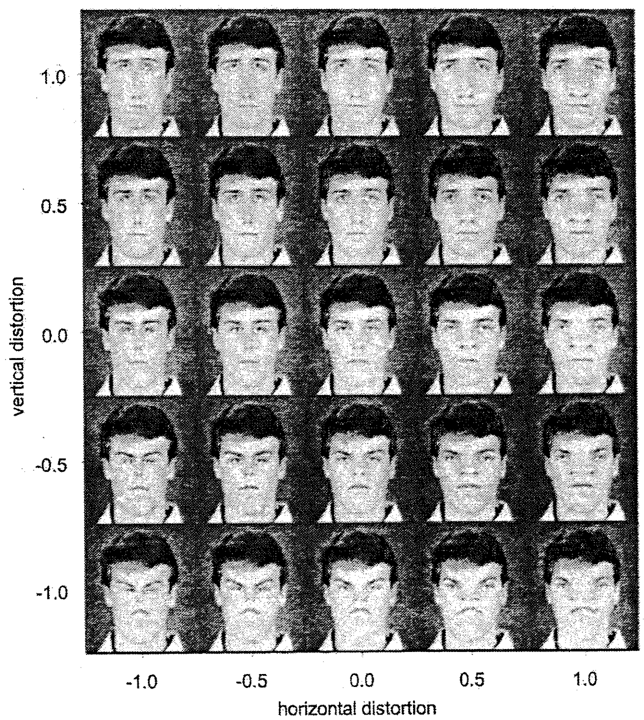

Fig. 1 An example of the array of face images, formed by graded expansions or contractions along the horizontal or vertical axes of the original image, shown at center.

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

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

The signs of w_x and w_y changed for pixels to either side of the midpoint. The amplitudes of the vertical and horizontal distortions (α_x and α_y) were varied independently over a range of -1 to $+1$ in steps of 0.02. In this way we generated a large (51×51) array of finely graded distortions, with the original image at the center. An example of the array is illustrated in Fig. 1, which shows the subset of images with α_x and α_y varied from -1 to $+1$ in steps of 0.5. In order to compare different adapting conditions, most measurements were based on the single face array shown in Fig. 1. A similar pattern of adaptation effects was confirmed on additional faces from the Matsumoto and Ekman set. Moreover, we have also found that there is a strong transfer of the adaptation effects across the different faces.[9]

The images were displayed on a standard color monitor controlled by a PC. Observers viewed the monitor binocularly in an otherwise dark room from a distance of approximately 1 m. At this distance the images subtended $8° \times 6.5°$. The images were centered on a $14° \times 18°$ uniform gray background of 30 cd/m^2 (comparable to the mean luminance of the images). The observers were told to view the images throughout but were not given specific fixation instructions.

A single experimental session lasted 1.5–2 h. During this session we measured the subset of images from the array that appeared "normal" to the observer (e.g., as a plausible image of a real face), either under neutral adaptation or after adaptation to a single distorted image. The

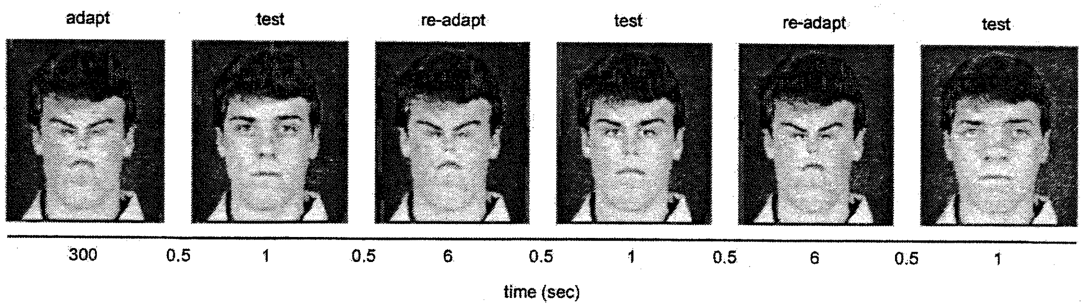

Fig. 2 Stimulus sequence for adaptation and testing. Test images varied on each trial to span different directions within the array, and were presented interleaved with a fixed adapting image.

measurements were made by using a staircase procedure to map out the boundary within the array defining the gamut of normal images. On each trial a single image was displayed for 1 s, and the observer responded whether the image was "normal" or "distorted" by pressing designated buttons on the keyboard. The image was drawn from one of eight staircases that varied the distortions along eight different directions at 45° intervals within the face array. A distorted response caused the next image from that staircase to shift closer to the array midpoint, while a normal response increased the distortion for the next trial. The array midpoint for the staircases was initially allowed to drift according to the observer's responses and was recalculated from the average of the reversal points at intervals of eight trials (during which images on each staircase were presented in random order). After 100 trials the midpoint was fixed for final estimates based on 8 reversals along each staircase.

Adaptation effects were assessed by repeating the rating task while test stimuli were interleaved with an adapting image. The adapting images were drawn from the extremes of the array and thus were highly distorted. Observers first viewed the adapting image for 300 s. The test trials were then alternated with 6 s readaptation intervals, with 0.5 s uniform fields interposed between each test and adapt image. This sequence is illustrated in Fig. 2. The neutral-adaptation ratings followed the same temporal sequence, but with a uniform field substituted for the adapting image. Each observer participated in only one experimental session testing the ratings for a single image and adapting distortion. No prior training or feedback was given. Measurements are reported for a total of 79 observers, who were undergraduate students at the University of Nevada and participated for course credit. The subjects had normal or corrected-to-normal acuity and were naive with respect to the purpose of the experiment.

3 Results

3.1 *Adaptation to Upright Faces*

Figure 3 shows an example of the normal versus distorted ratings measured for six different individual observers. The symbols plot the mean reversal points for the eight staircases, either under neutral adaptation to the uniform screen (open circles), or following adaptation to a distorted face (closed circles). In the top row, results are plotted for observers who adapted to a face that was fully expanded along both the vertical and horizontal axes ($\alpha_x = +1$, $\alpha_y = +1$), while the bottom row shows settings when the adapting face was instead fully contracted ($\alpha_x = -1$, $\alpha_y = -1$). Individual differences in the ratings are substantial, suggesting that observers adopted very different criteria in judging the images. Nevertheless, the pre-adapt settings (open circles) define a circumscribed region near the center of the array, and adaptation to the distorted images shifts these settings in consistent ways. In particular, following adaptation the reversal points are clearly biased toward the adapting image. For some observers, these biases were so large that the adapting image itself came to be rated as normal, while the original image came to be judged as clearly distorted. These after-effects in the original face are consistent with the after-effects we measured previously, and again suggest that adaptation to a distorted face causes the original face to appear distorted in the opposite direction. Moreover, the results are consistent with the suggestion that these after-effects result from a perceptual renormalization of the face array, so that the adapting image appears more normal or neutral.[9] (It should be noted that many subjects spontaneously commented on this renormalization. When the distorted adapting faces were first shown they appeared highly disfigured and often evoked comments to this effect or other overt responses such as laughter. Yet a number of subjects commented that with continued viewing the adapting face looked progressively "better," consistent with their ratings.)

The solid lines in Fig. 3 plot ellipses estimated from a least-squares fit to the reversals. These ellipses in general provided a good account of the observed reversal points, and thus provided a succinct description of the area within the array that was perceived as normal. In subsequent figures we have therefore used these ellipses to represent the ratings. To extract general features of the ratings, we further calculated the ellipses based on the reversal points averaged across observers. In Fig. 4, these averaged ellipses (each based on three observers) are shown for nine different adapting conditions, including the original face or one of eight different distortions drawn from the extremes of the array. Again, the ratings tend to be centered around the original face before adaptation, but are systematically shifted toward each adapting image following adaptation. The shifts in the mean of the ratings are comparable in magnitude to those measured with a recognition task.[9] (Interestingly, Fig. 4 suggests that after-effects may be stron-

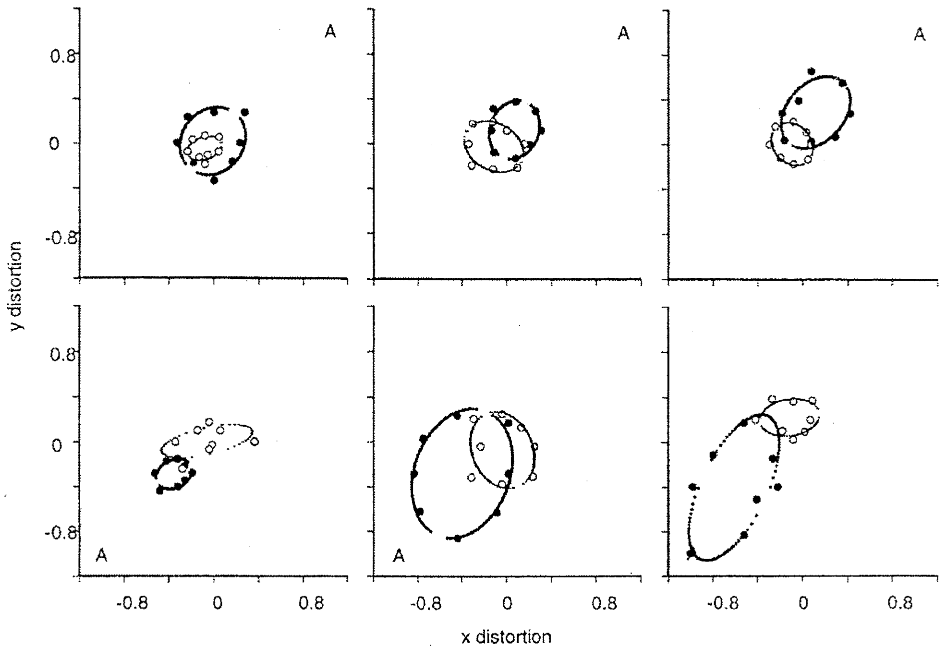

Fig. 3 Examples of the "normal vs distorted" ratings for images in the face array. Each panel plots the staircase reversal points for the ratings made by a single observer before (open circles) or after (closed circles) adaptation to a distorted face image (at the array location indicated by the symbol "A"). Ellipses fitted to the reversals define the boundary of images in the array that appeared normal.

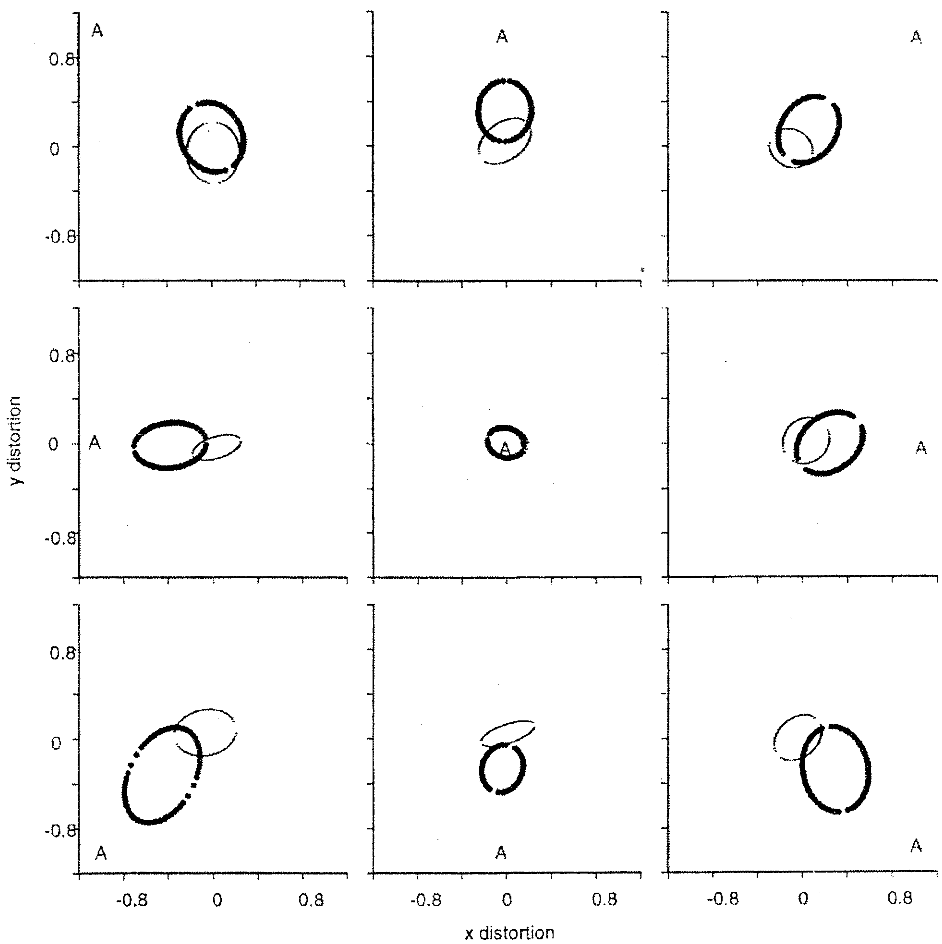

Fig. 4 Mean ratings before (thin ellipse) or after (thick ellipse) adaptation to the original face or eight different distorted faces. Each panel plots the ratings for a single adapting image (at a location in the array indicated by the symbol "A").

ger for contracted adapting images, a difference we also observed in our earlier measurements.[9])

An exception to these changes is seen in the results for adaptation to the original image. As shown in the center panel of Fig. 4, in this condition the pre- and postadapt ellipse remain very similar (and in fact cannot be separately resolved in Fig. 4). The absence of an after-effect when adapting to the original image is consistent with the asymmetries in the after-effect reported by us previously.[9] In that work we found that while adaptation to a distorted image had a large effect on the appearance of the original image, adapting to the original image conversely induced little change in the appearance of the distorted image. These results are of interest because they suggest that adaptation is *not* altering perception by causing other images to appear less like the adapting image. Such biases should push the perception of all images away from the array coordinates of the adaptation, and thus for adaptation to the original image should shrink the perceived gamut of normal images toward the adaptation point. Figural after-effects of this form are common, and gave rise to the notion of the "distance paradox" in after-effects (in which all stimuli appear biased away from the adapting stimulus, while the adaptation pattern itself appears unaltered[16,17]). For example, adapting to a 2 c/deg grating causes frequencies above 2 c/deg to appear higher in frequency than they did before adaptation, and gratings below 2 c/deg to appear lower, but does not alter the perceived frequency of the adapting grating itself.[4] Thus a fixed perceptual range (e.g., one octave centered on the adapting frequency) should correspond to a smaller physical range (i.e., less than an octave difference in the actual frequencies). The distance paradox occurs for stimulus dimensions that are represented as a continuum along which no particular value is special. That adaptation to an undistorted face image does not change judgments of the distortions suggests instead that the face distortions may not be represented as a continuum, because stimuli like the original have a neutral value, analogous to the status of "white" in color perception or "stationary" in motion perception. In such cases, a prominent component of adaptation after-effects may be a perceptual "recentering" of the array, so that the prevailing image appears more neutral. As noted above, this is also consistent with our observations that the adapting stimulus itself appears less distorted with prolonged viewing. (An alternative, which we discuss below, is that there may be no apparent after-effects for adaptation to the original face because observers came into the experiment already under a state of adaptation to a similar "average" face.)

However, the results for the distorted adapting images also reveal that adaptation is not merely inducing a mean shift in observers' judgments of faces within the array, for the *range* of the settings is also changed in systematic ways. Specifically, Fig. 4 suggests that the area rated as normal increases under biased states of adaptation. This change is further illustrated in Fig. 5, which shows a scatter plot of the areas of individual observers' ellipses measured before or after adaptation. Systematic increases in the range of normal ratings are evident, implying that observers either become more tolerant or less sensitive to the image distortions following the adaptation. A second feature suggested by the results is that these changes in area are selective, for there is a tendency for the ellipses to become oriented along the direction of the adapting distortion. For example, in Fig. 3 the settings for the six observers do not show a consistent orientation before adaptation, but are consistently elongated along the positive diagonal following adaptation. This suggests that the adaptation caused observers to become selectively more tolerant or less sensitive to the axes defining the adapting distortion.

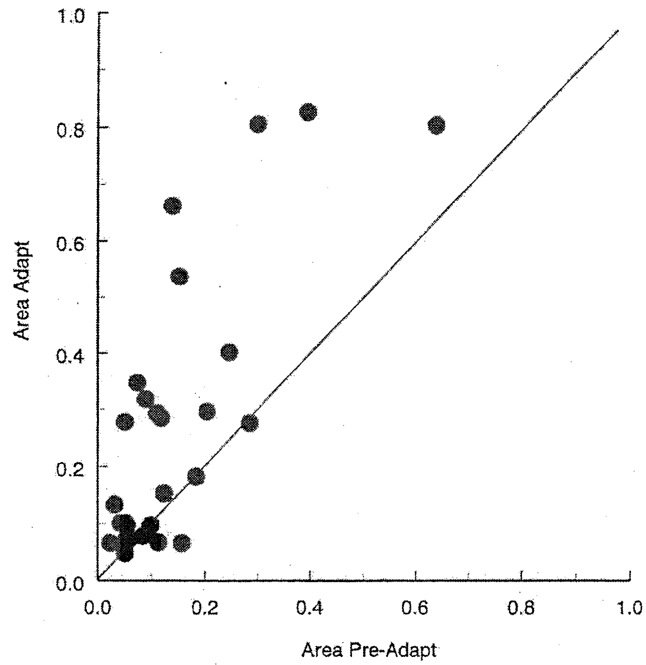

Fig. 5 Changes in the range of images perceived as normal after adaptation. The area of the ellipses fitted to individual observers' reversal points are plotted before or after adaptation.

3.2 *Time Course of the Adaptation*

Large after-effects in the appearance of faces can be induced by only brief periods of adaptation.[9] For example, we have found in informal observations that exposure times of well under 1 min are sufficient to induce salient biases. This suggests that adaptation could adjust quickly enough to often play a significant role in modulating perception under normal viewing conditions. However, the time constants governing the face adaptation, or figural after-effects in general, are not well defined. We chose a testing procedure that was designed to frequently readapt the observer in order to maintain the adaptation during the course of the ratings. Yet despite the initial 5 min exposure, drifts in the mean ratings (toward larger after-effects) were often evident over the first 100 trials during which the mean was allowed to vary. It is thus likely that the actual after-effects are underestimated by our measurements. For a separate set of subjects we explored the buildup of the after-effect by measuring the ratings during the course of 500 trials. Each trial was again preceded by a 6 s top up but the runs in this case began without the initial 300 s adapt period. Under these conditions subjects showed a steady and progressive buildup of the after-effect. This indicates that the biases are indeed the result of adaptation and not simply of a "contrast effect" between the test and the immediately preced-

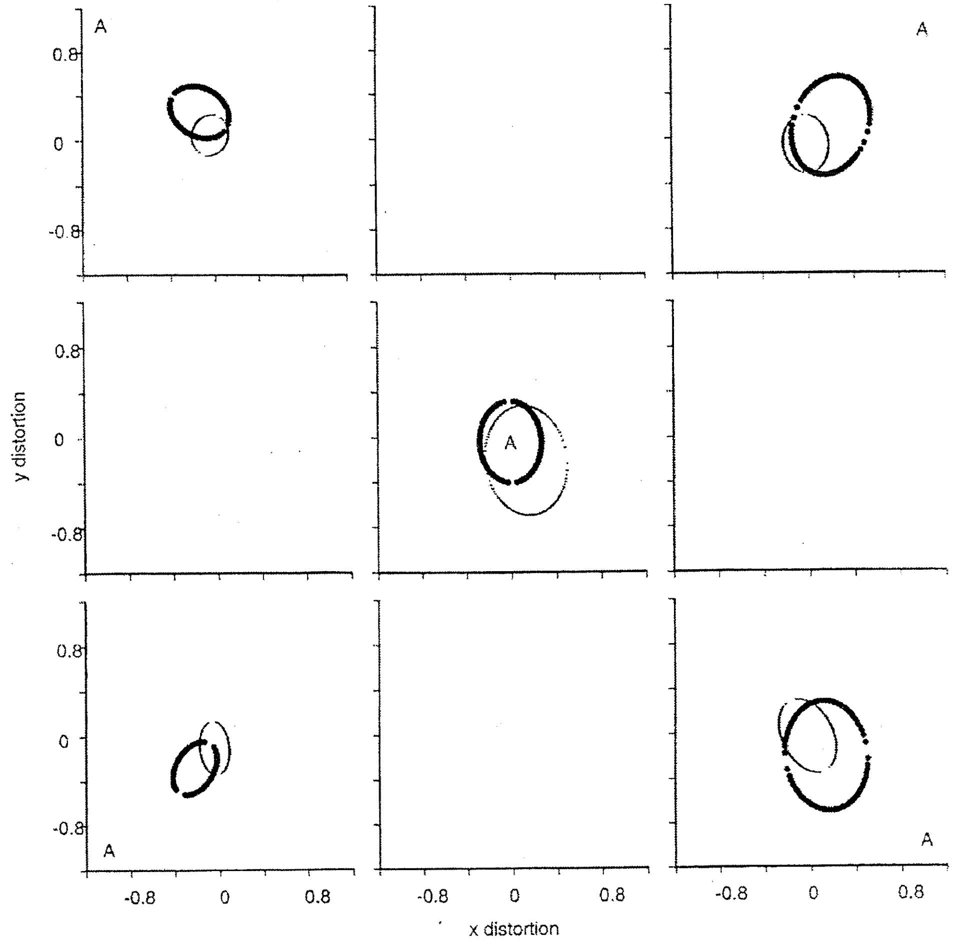

Fig. 6 "Normal vs distorted" ratings for inverted test and adapting images. The panels plot the ellipses fitted to the average ratings of three observers either before (thin line) or after (thick line) adaptation to a single distortion (indicated by the symbol "A").

ing image. Moreover, for some subjects the settings had not clearly asymptoted even at the end of the trials (i.e., after more than 1 h). Our results therefore do not preclude the possibility that with sufficiently long exposures observers could adapt (renormalize) completely even to the large image distortions we tested. We have not measured the decay of the after-effect, yet we have again found in informal observations that when the test face is viewed continuously it appears to slowly return to the original configuration over the course of several seconds. These changes are salient and rapid enough that the continuously viewed test image can appear animated, although changes in the appearance of the test image were not evident during the 1 s presentations used in the rating task.

3.3 *Adaptation to Inverted Faces*

The preceding results show that judgments of distortions in faces can be strongly influenced by changes in the observer's state of adaptation. To what extent might these adaptational adjustments reflect effects that are characteristic of face perception? As noted above, a number of sources of evidence suggest that face perception depends on specialized processes that may be carried out within specialized areas of the cortex.[18–20] Psychophysically, a principal argument for these processes has come from the finding that faces become much more difficult to recognize when the images are presented upside down,[21,22] and distortions in the faces become much less noticeable.[23] We therefore examined the effects of inverting the images on judgments of the image distortions and how they change with adaptation.

Figure 6 shows the ratings when both the test and the adapting images were inverted. In this case we measured the after-effects for five adapting conditions (the original face and images at the four corners of the array). Each panel again shows the preadapt and postadapt settings averaged across three observers. There are once again systematic biases in the ratings following adaptation, so that the images judged as normal shift toward each adapting distortion. These biases are similar in magnitude for the upright and inverted images, and comparable effects for upright and inverted images were also found when the adaptation was measured with a recognition task.[9] This suggests that the potency of a distortion as an adapting stimulus might not depend on how sensitive observers are to the distortion. However, we cannot in fact be certain that observers were less sensitive to the distortions in our inverted stimuli. The average area of the preadapt ellipses was almost twice as large for the inverted faces compared to the upright faces, yet this difference was not significant ($t=1.6$, $df=50$; $p>0.1$).

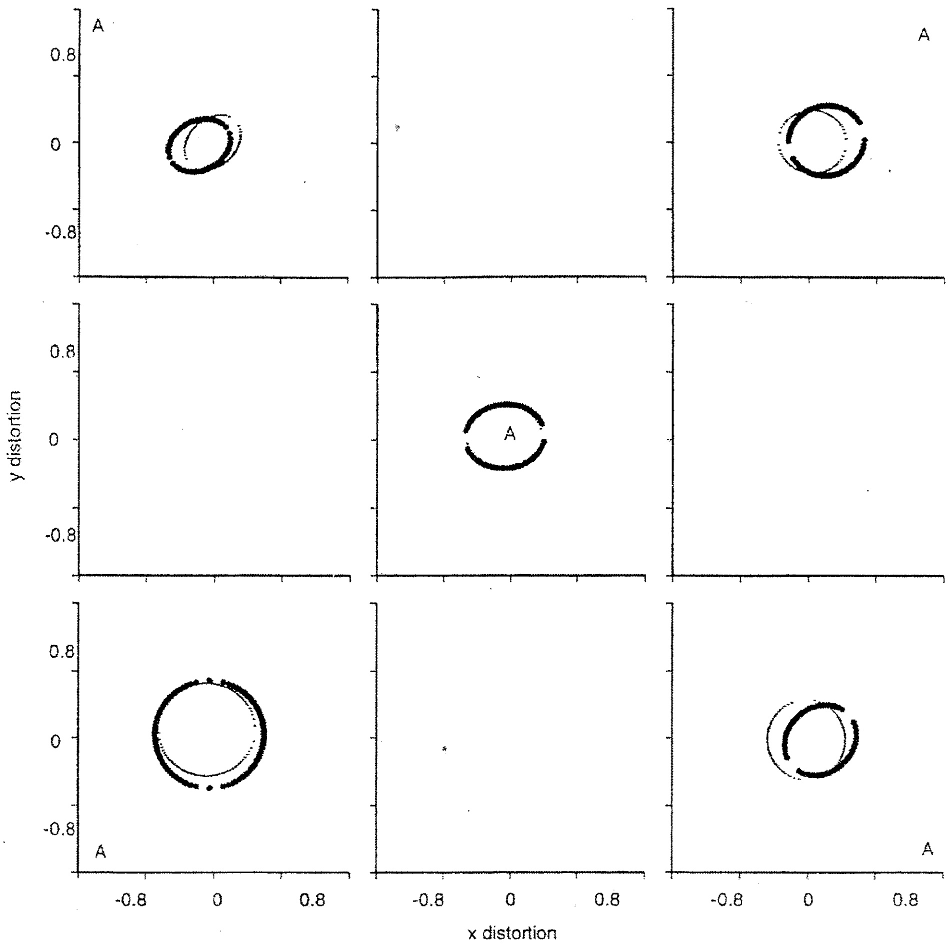

Fig. 7 "Normal vs distorted" ratings for upright test images and inverted adapting images. The panels plot the ellipses fitted to the average ratings (for three observers) before (thin line) or after (thick line) adaptation to a single distortion (indicated by the symbol "A").

In Fig. 7, results are plotted for cases in which the adapting image was inverted while the test image was upright. The after-effects are now clearly weaker and there are at best only modest suggestions of a change in the ratings with adaptation. Shifts in the means of the ratings are slight and comparable to the marginal after-effects measured by measures of face recognition.[9] Moreover, the present ratings show that there is little change in the range of stimuli perceived as normal, and little bias in the orientation of this range. Note that inverting the images changes the structure of the image, but does not alter the image distortion, for the distortion is symmetric about the vertical and horizontal axes. The absence of after-effects for these conditions thus indicates that the after-effects are selective for the structure of the distorted image.

3.4 *Perceived Distortions in Nonface Images*

As we noted, we examined adaptation effects in images of faces because changes in facial configuration are highly salient. Moreover, the range of configurations over which faces can normally vary is limited, so that the "permissible" distortions in the images should be clearly constrained. For many other classes of natural images these constraints may be much weaker, because the images themselves may naturally vary more widely, observers may have less *a priori* knowledge or expectations about these variations, and the different elements of the image may be less closely coupled as part of a common configuration. To assess the impact of these factors, we compared judgments of distortions when the same physical distortions used to construct the face arrays were applied to images of outdoor scenes.

Figure 8 shows an example of the image arrays for two images we tested, a distant hillside or a closeup of tree bark. It is evident from inspection of Fig. 8 that the distortions are less apparent, and that many of the array images, even though they represent the extremes of the distortions we examined, still appear as plausible or normal images. Individual ratings of the distortions are shown in Fig. 9 for six observers, for the landscape (top row) or bark (bottom row). There are again interobserver differences, yet in general the gamut of normal images is several times larger than the average gamut for the face images (replotted as the thin ellipse in each panel). Moreover, for many observers the gamut extended to the limits of the image array, so that we could not in fact define the boundary of the images rated as normal. We also tested whether adaptation could bias the ratings for these images, but could not confirm a systematic

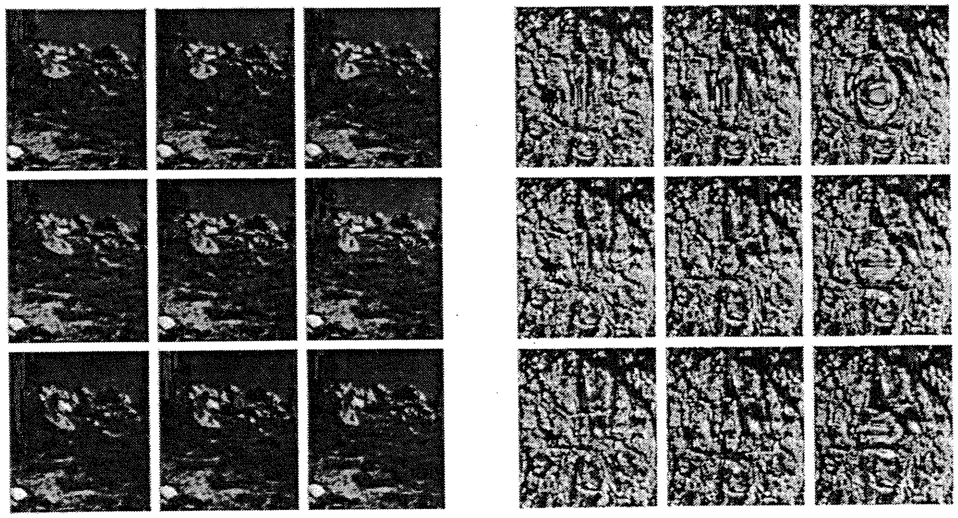

Fig. 8 Image arrays for two natural images, formed by applying the same distortions used to create the face arrays. The panels show the largest distortions in the array (corresponding to α_x and α_y values of -1, 0, or $+1$, with the original, undistorted images at the center).

figural after-effect for these stimuli, despite the fact that adaptation to the same images does induce large and systematic changes in threshold contrast sensitivity.[13,24] These results thus illustrate that sensitivity to the same physical distortions—and perhaps also the adaptational influences on the perceived distortions—may vary widely for different classes of images.

4 Discussion

We have shown that prior exposure to distorted images can produce enormous changes in observers' perceptions and ratings of image distortions. Thus variations in the state of adaptation can have a profound influence on judgments of images, greatly increasing the range of stimuli that could

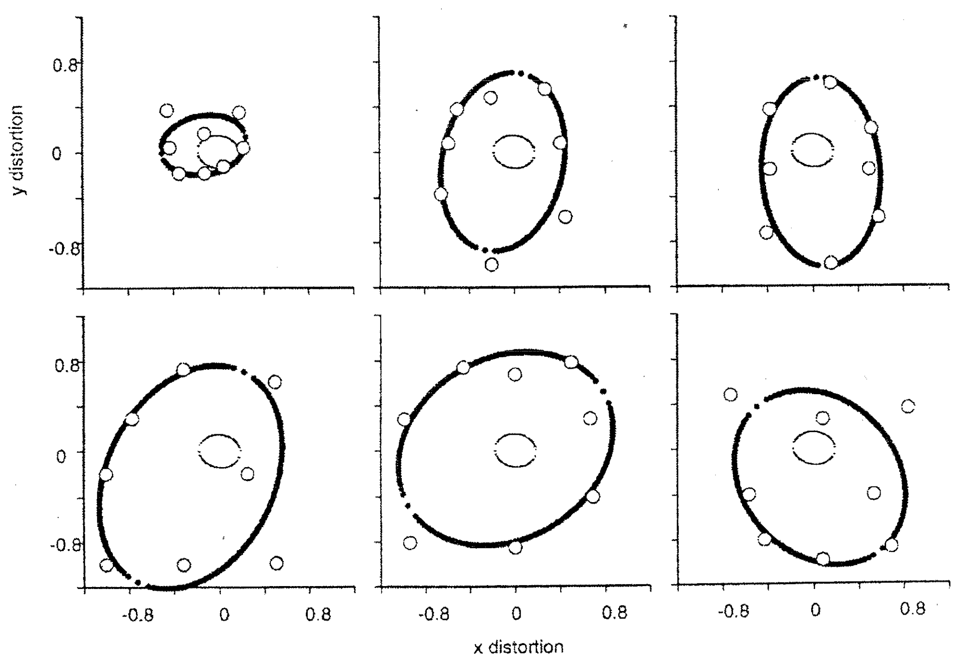

Fig. 9 Preadapt ratings of the distortions in the landscape (top row) or bark (bottom row) images. Each panel shows the settings for a different observer. Open circles plot the staircase reversal points for the ratings. Thick lines plot the best fitting ellipse. The preadapt ellipse for the face image (thin line) is plotted for comparison.

appear normal or distorted. That these large after-effects occur after freely viewing images for only brief periods suggests that changes in adaptation may often play an important role in influencing observers' perception and judgments within everyday contexts.

We compared the adaptation effects for biased images with the ratings obtained in "preadapt" or "neutral adaptation" control runs. However, such labels are in some sense misleading, for it is more reasonable to suppose that the visual system is always under some state of pattern adaptation, and that these states determine the properties of pattern perception in the same way that light adaptation plays a ubiquitous role in defining sensitivity. Thus our preadapt ratings may in reality depend on the adaptation states that observers bring into the experiment, and it is notable in this regard that adaptation to the original face did not substantially alter observers' judgments. Interobserver differences in the preadapt settings were large, and surprising because these settings should in part reflect observers' *a priori* expectations about normal variations in the population of faces. Moreover, different directions within the array alter the face in different ways that may be more or less plausible in the natural population. For example, distortions along the positive diagonal produce an overall expansion or contraction of the face, while the negative diagonal instead alters the aspect ratio of the facial configuration, yet our results provide no evidence that observers as a group are more tolerant to one source of variation than the other. We intentionally chose to give minimal instructions for the ratings to allow subjects to make judgments that were natural to them, and thus the differences across observers may reflect differences in criteria as much as differences in expectations. It remains an intriguing question to ask to what extent observers' expectations for image variations are matched to the natural variations in different object classes. For example, a number of studies have characterized physical patterns of variation among faces (see, e.g., Refs. 14, and 25–28). An interesting extension of the present results would be to test the normal versus distorted judgments along these stimulus axes to examine how these judgments are related to the actual distribution of faces.

Despite interobserver differences in the ratings, adaptation to biased images produces clear and consistent biases in the appearance of the images, and in the present study we have shown that this can markedly alter the range of face stimuli that observers rate as normal. Because these adaptation effects reflect changes in visual perception, they should similarly affect other kinds of judgments about faces. For example, a number of studies have found that average faces are rated as more attractive.[15,29,30] However, our results suggest that adaptation can strongly bias the appearance of an average face, and that because of adaptation distinctive faces should appear more normal or average the longer they are viewed. Thus judgments of such qualities as attractiveness and distinctiveness may strongly depend on the specific, recent history of faces to which individual observers have been exposed and to which they are therefore currently adapted. [An interesting anecdotal observation in line with this prediction was reported by Malinowski, who noted that in the course of living in the Trobriand Islands his judgments of beauty began to agree with the Trobrianders' judgments (see Ref. 31, p. 196). It is tempting to speculate that these changes reflected actual shifts in visual perception rather than "criteria," and that Malinowski was simply adapting perceptually to his new visual environment.]

We earlier noted[9] that the tendency for adaptation to renormalize face perception is functionally analogous to the effects of chromatic adaptation, which recalibrates sensitivity so that the prevailing color signal appears white. For example, biasing an illuminant toward longer wavelengths induces selective changes in chromatic sensitivity which tend to discount the illuminant change, so that the biased illuminant appears neutral. These compensatory adjustments produce corresponding changes in all other colors, resulting in large changes in the mean perceived color of stimuli. In color vision these mean changes primarily reflect sensitivity adjustments in early cone-specific pathways,[8] but analogous shifts in the neutral point are found in some cortical adaptation effects such as the motion after-effect.[32] Similarly, the after-effects of face adaptation appear to reflect changes in the face stimulus that appears neutral, inducing corresponding shifts in the appearance of other faces in the array.

However, the present results suggest that the adaptation not only realigns perception but rescales it as well, for the increases in the range of the ratings under adaptation imply that under biased adaptation states observers have less sensitivity to variations around the neutral point. One possible basis for this effect is that sensitivity to the distortions is not uniform throughout the array, or that the realignment is not simply a linear translation in the responses to the stimuli. For example, in chromatic adaptation the mean shifts in color appearance result from a multiplicative rescaling of the cone signals. Thus, for example, increasing the average signal in short-wavelength sensitive cones reduces their gain. This not only shifts the white point (toward higher *S*-cone excitation), but also decreases sensitivity to linear changes in *S*-cone excitation around the white point. Sensitivity or adaptational shifts to the distortion gradients we examined might similarly be nonlinear. A second possibility is that the after-effects reflect more than one type of sensitivity adjustment. To draw again from color perception, the visual system adjusts not only to the average color in scenes through chromatic adaptation, but also to the variations in color around the average, through chromatic contrast adaptation.[8] These latter adjustments result in selective losses in sensitivity to the color directions defining the chromatic contrast. The adapting face distortions we examined might similarly represent "contrasts" in face stimuli, and adaptation to these contrasts might alter sensitivity to these contrasts. A more direct test of this form of adaptation was proposed by Atick, Griffin, and Redlich,[25] who suggested adapting to face images that are temporally modulated along different gradients.

Acknowledgments

Supported by NEI Grant No. EY-10834. The authors thank Seth Roberts for pointing out the reference to Malinowski. A disk with programs demonstrating face adaptation is available on request from author M.A.W.

References

1. A. Wohlgemuth, "On the after-effect of seen movement," *Br. J. Psychol. Monogr. Suppl. No.* **1**, 1–117 (1911).
2. J. J. Gibson and M. Radner, "Adaptation, after-effect and contrast in the perception of tilted lines. I. Quantitative studies," *J. Exp. Psychol.* **20**, 453–467 (1937).
3. W. Kohler and H. Wallach, "Figural aftereffects: An investigation of visual processes," *Proc. Am. Philos. Soc.* **88**, 269–357 (1944).
4. C. Blakemore and P. Sutton, "Size adaptation: A new aftereffect," *Science* **166**, 245–247 (1969).
5. R. B. Welch, "Adaptation of space perception," in *Handbook of Perception and Human Performance*, K. R. Boff, L. Kaufman, and J. P. Thomas, Eds., p. 24.1–24.45, Wiley, New York (1986).
6. J. J. Gibson, *The Ecological Approach to Visual Perception*, Houghton Mifflin, Boston (1979).
7. J. Addams, "An account of a peculiar optical phenomenon seen after having looked at a moving body, etc.," *London, Edinburgh, Dublin Philos. Mag. J. Sci.* **5**, 373–374 (1834).
8. M. A. Webster, "Human colour perception and its adaptation," *Network Comput. Neural Syst.* **7**, 587–634 (1996).
9. M. A. Webster and O. H. MacLin, "Figural after-effects in the perception of faces," *Psychonomic Bull. Rev.* **6**, 647–653 (1999).
10. M. J. Farah, K. D. Wilson, M. Drain, and J. N. Tanaka, "What is 'special' about face perception?," *Psychol. Rev.* **105**, 482–498 (1998).
11. M. A. Webster and O. MacLin, "Visual adaptation and the perception of distortions in natural images," in *Human Vision and Electronic Imaging III*, B. E. Rogowitz and T. N. Pappas, Eds., *Proc. SPIE* **3299**, 264–273 (1998).
12. D. Matsumoto and P. Ekman, "Japanese and caucasian facial expressions of emotion (JACFEE) and neutral faces (JACNeuF)," Department of Psychology, San Francisco State University, San Francisco, CA (1988).
13. M. A. Webster and E. Miyahara, "Contrast adaptation and the spatial structure of natural images," *J. Opt. Soc. Am. A* **14**, 2355–2366 (1997).
14. A. J. O'Toole, T. Vetter, H. Volz, and E. M. Salzer, "Three-dimensional caricatures of human heads: Distinctiveness and the perception of facial age," *Perception* **26**, 719–732 (1997).
15. G. Rhodes and T. Tremewan, "Averageness, exaggeration, and facial attractiveness," *Psychol. Sci.* **7**, 105–110 (1996).
16. C. E. Osgood and A. W. Heyer, "A new interpretation of figural after-effects," *Psychol. Rev.* **59**, 98–118 (1952).
17. G. von Bekesy, *Experiments in Hearing*, McGraw–Hill, New York (1960).
18. R. Desimone, "Face-selective cells in the temporal cortex of monkeys," *J. Cognitive Neurosci.* **3**, 1–8 (1991).
19. M. J. Farah, K. D. Wilson, H. M. Drain, and J. R. Tanaka, "The inverted face inversion effect in prosopagnosia: Evidence for mandatory, face-specific perceptual mechanisms," *Vision Res.* **35**, 2089–2093 (1995).
20. N. Kanwisher, J. McDermott, and M. M. Chun, "The fusiform face area: A module in human extrastriate cortex specialized for face perception," *J. Neurosci.* **17**, 4302 (1997).
21. R. K. Yin, "Looking at upside down faces," *J. Exp. Psychol.* **81**, 141–145 (1969).
22. T. Valentine, "Upside-down faces: A review of the effect of inversion upon face recognition," *Br. J. Psychol.* **79**, 471–491 (1988).
23. P. Thompson, "Margaret Thatcher: A new illusion," *Perception* **9**, 483–484 (1980).
24. M. A. Webster, "Contrast sensitivity under natural states of adaptation," in *Human Vision and Electronic Imaging IV*, B. E. Rogowitz and T. N. Pappas, Eds., *Proc. SPIE* **3644**, 58–70 (1999).
25. J. J. Atick, P. A. Griffin, and A. N. Redlich, "The vocabulary of shape: Principal shapes for probing perception and neural response," *Network Comput. Neural Syst.* **7**, 1–5 (1996).
26. A. M. Burton, V. Bruce, and M. Dench, "What's the difference between men and women? Evidence from facial measurement," *Perception* **22**, 153–176 (1993).
27. P. W. Hallinan, G. G. Gordon, A. L. Yuille, P. Giblin, and D. Mumford, *Two- and Three-Dimensional Patterns of the Face*, A. K. Peters, Ed., Natick, MA (1999).
28. M. Turk and A. Pentland, "Eigenfaces for recognition," *J. Cognitive Neurosci.* **3**, 71–86 (1991).
29. K. A. Deffenbacher, T. Vetter, J. Johanson, and A. J. O'Toole, "Facial aging, attractiveness, and distinctiveness," *Perception* **27**, 1233–1243 (1998).
30. J. H. Langlois and L. A. Roggman, "Attractive faces are only average," *Psychol. Sci.* **1**, 115–121 (1990).
31. D. Symons, *The Evolution of Human Sexuality*, Oxford University Press, New York (1979).
32. *The Motion Aftereffect*, G. Mather, F. Verstraten, and S. Anstis, Eds., MIT Press, Cambridge, MA (1998).