Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Vision Research 44 (2004) 113–117
Vision Research
www.elsevier.com/locate/visres

Rapid Communication

Depth-dependent blur adaptation

Peter W. Battaglia [a], Robert A. Jacobs [b,*], Richard N. Aslin [b]

[a] *Department of Psychology, University of Minnesota, Minneapolis, MN 55455, USA*
[b] *Department of Brain and Cognitive Sciences and the Center for Visual Science, University of Rochester, 274 Meliora Hall, Rochester, NY 14627-0270, USA*

Received 4 July 2003; received in revised form 3 September 2003

Abstract

Variations in blur are present in retinal images of scenes containing objects at multiple depth planes. Here we examine whether neural representations of image blur can be recalibrated as a function of depth. Participants were exposed to textured images whose blur changed with depth in a novel manner. For one group of participants, image blur increased as the images moved closer; for the other group, blur increased as the images moved away. A comparison of post-test versus pre-test performances on a blur-matching task at near and far test positions revealed that both groups of participants showed significant experience-dependent recalibration of the relationship between depth and blur. These results demonstrate that blur adaptation is conditioned by 3D viewing contexts.

Keywords: Perceptual learning; Visual adaptation; Perceptual constancy; Visual blur

1. Introduction

Blur is an inherent property of retinal images that contain two or more objects lying in two or more depth planes. Despite the presence of blur in most retinal images, particularly of natural scenes, the visual world typically appears 'in focus'. Recent data suggest that observers' perceptions of image blur adapt over time based on the blur of recently viewed images. Webster, Georgeson, and Webster (2002) showed that observers who were exposed to blurry images for a few minutes tended to perceive normal images as overly sharp, whereas observers who were exposed to sharp images perceived normal images as blurry. This finding suggests that our neural representations are adjusted so that image information at different spatial frequency scales fall within the same limited dynamic range. In other words, these adjustments allow our neural mechanisms to operate in a manner that is roughly invariant to shifts in the distribution of contrast across different spatial scales. If so, then adaptation to image blur can be regarded as a type of perceptual constancy.

We examined the hypothesis that neural adjustments to image blur are sensitive to visual depth. There are several reasons for us to expect that this is the case. For example, consider an observer viewing a scene with multiple objects. When fixating an individual object at a close viewing distance, the focal plane of each of the observer's eyes is adjusted to bring the object into focus on the two retinas by accommodation of the lens. The images of objects far from the focal plane are, however, blurry. Nonetheless, these objects often appear to the observer to be in focus. We hypothesize that this is due to neural adaptation. Although the image of an object far from the focal plane is blurry, the image of this object could easily be brought into focus through re-accommodation of the lens. Consequently, the observer's brain "knows" that the blurriness of the image of the object far from the focal plane is a temporary feature due to current viewing conditions, and not an intrinsic property of the object. We conjecture that it also knows how to adapt so as to "de-blur" the image of this object based on its distance from the focal plane. The experiment reported below demonstrates that the relationship between the neural de-blurring process and visual depth is adaptable as a function of training experience. We conclude, therefore, that neural adaptation to image blur is sensitive to visual depth.

* Corresponding author. Tel.: +1-585-275-0753; fax: +1-585-442-9216.
E-mail address: robbie@bcs.rochester.edu (R.A. Jacobs).

0042-6989/$ - see front matter
doi:10.1016/j.visres.2003.09.010

2. Methods

2.1. Participants

Eight adults (students at the University of Rochester between ages 18 and 25 years of age) with normal or corrected-to-normal acuity and normal stereopsis were paid $8 each for participation in a one-hour experimental session. Participants were treated in accordance with the ethical guidelines established by the University's Research Subjects Review Board.

2.2. Stimuli and apparatus

Participants placed their heads on a chinrest located 23 cm from a 19 in Sony Trinitron monitor. They viewed the monitor while wearing Crystal Eyes wireless stereoscopic glasses that alternated video frames between the eyes at 50 Hz. An interocular separation of 6.5 cm was used to calculate retinal disparities.

Each stimulus display consisted of a fixation target and two fronto-parallel surfaces (see Fig. 1). The fixation target consisted of a diamond-shaped surface covered with a red checkerboard pattern and a purple cross. The cross hovered directly over the surface and rotated slowly. This fixation target was located in the center of the display, and its distance from the observer along the depth axis was 32.9 cm. It was designed to provide a mildly interesting visual stimulus that would reduce the participant's desire to look directly at either of the two adjacent fronto-parallel surfaces. The two fronto-parallel surfaces were 19.1 cm tall and 12.5 cm wide. One surface was placed on each side of the fixation target, and the surfaces were separated along the horizontal axis by 3.3 cm. The fixation target and the fronto-parallel surfaces were rendered in depth using stereo and perspective cues.

On each trial of the experiment, an image of a grassy texture was placed on each fronto-parallel surface. This texture was filtered to produce 16 different magnitudes of image blur. The set of blurred images was created by convolving an unfiltered image of the texture with 16 first-order Butterworth filters that differed in their cutoff frequency. The Nyquist frequency of the unfiltered image, given the distance from the observer to the monitor, was 7.075 cycles per degree of visual angle (cyc/deg).

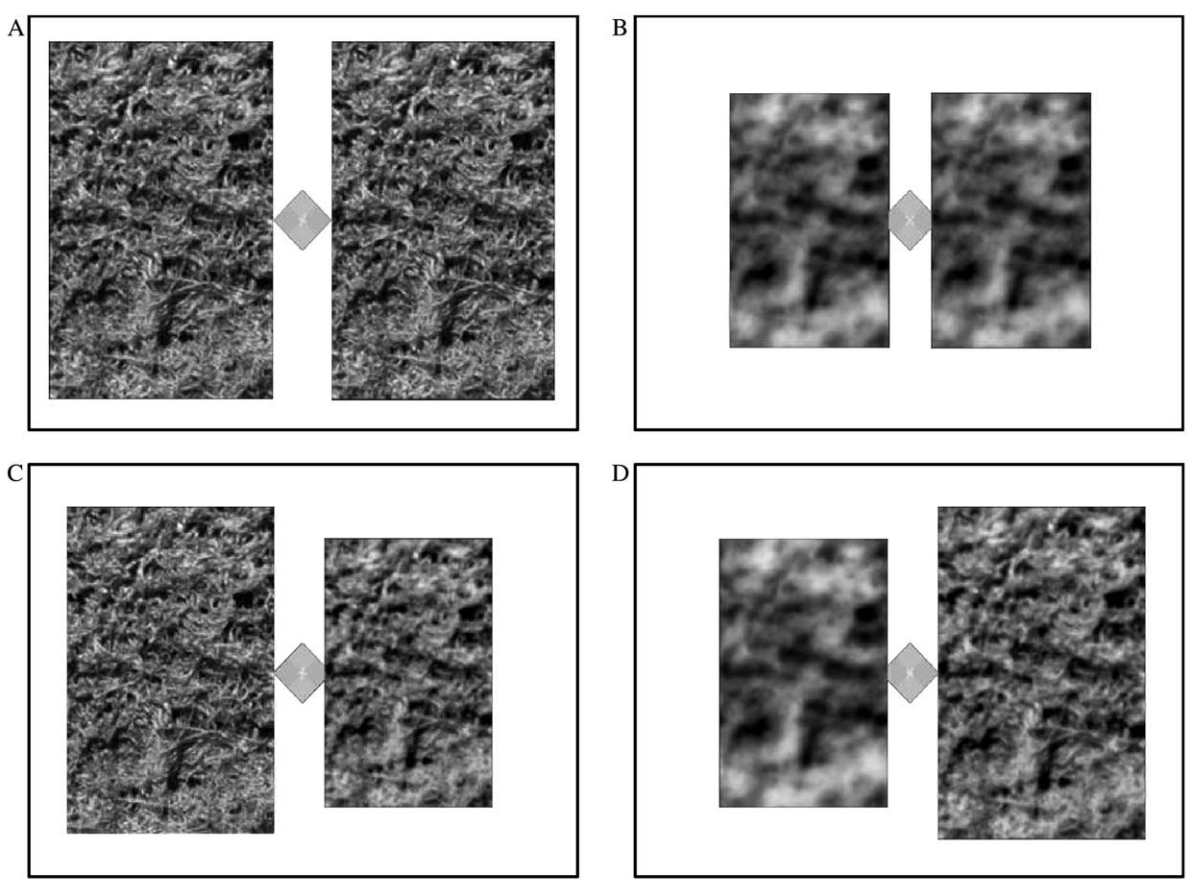

Fig. 1. Screen captures of the experimental stimuli. The fixation target is the diamond-shaped object between the two surfaces. (A) Adaptation phase: both surfaces are near and sharp. (B) Adaptation phase: both surfaces are far and blurry. (C) Test phase: the left surface is near and sharp; the right surface is far and of standard (middle-level) blur. (D) Test phase: the left surface is far and blurry; the right surface is near and of standard (middle-level) blur.

The lowest cutoff frequency we used was 0.142 cyc/deg and the highest was 3.11 cyc/deg. As a matter of notation, let the blur magnitudes used in the experiment be numbered from 0 (most blurry) to 15 (least blurry). These blur magnitudes can be converted into frequency cutoffs using the equation

$$\mathrm{sf} = 0.142\left(\frac{7.075}{0.142}\right)^{b/19}$$

where sf is the spatial frequency cutoff and b is the blur magnitude. Differences between images with successive blur magnitudes were roughly equally easy for participants to perceive. Although filtering did not affect the mean luminances of the images, it did lower their RMS pixel contrasts. After filtering, therefore, each image's RMS pixel contrast was restored to that of the unfiltered image by linearly scaling the image's pixel luminance values.

2.3. Procedure

The experiment contained four stages: baseline training, pre-test, adaptation training, and post-test. During Stage 1, baseline training, each participant fixated the central fixation target for 3 min 15 s while the fronto-parallel surfaces traveled to and from the participant along the midline depth axis, covering a range from 23 to 32.2 cm (this range is known as the adaptation range; see Fig. 2). The speed of the surfaces was set so that it took approximately 10 s for the surfaces to move from the point furthest from the participant to the point closest to the participant and then back to the furthest point. The surfaces always traversed the same depth range from the participant in tandem, and had the same image mapped onto them (a mid-range blur magnitude equal to 8). Importantly, we did not vary this image as the surfaces traversed the adaptation range (i.e. to the participants, displays differed over time soley due to perspective projection of the surfaces and the images mapped onto them).

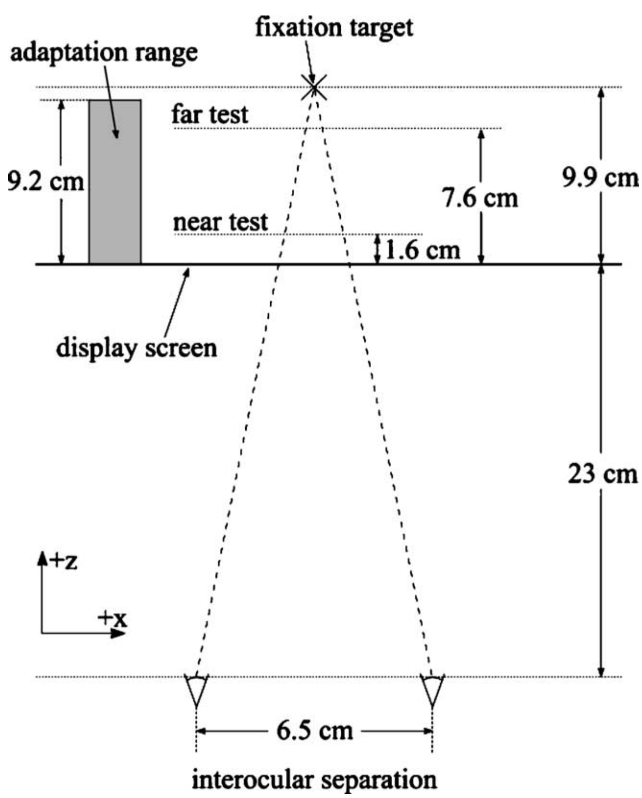

Fig. 2. Schematic of display scene presented to the participants.

Stage 2 (the pre-test) was used to measure participants' baseline blur-matching performance at both a near test position and a far test position. Consider for the moment just those test trials relevant to assessing a participant's performance at the near test position—the even-numbered trials. On each test trial, participants fixated the central fixation target and viewed the fronto-parallel surfaces for 500 ms. The surfaces were stationary; one was positioned 24.6 cm from the participant and the other was at a distance of 30.6 cm (the selection as to whether the left or right surface would be near or far was randomized). Participants viewed the images on the surfaces and judged whether the image on the left surface was more blurry or less blurry than the image on the right surface. The near surface contained the "standard" image, which had a blur magnitude of 8 (0.74 cyc/deg). The far surface contained the "comparison" image. The blur of the comparison image was adjusted on a trial-by-trial basis according to a 1-up/1-down (in units of blur magnitude) adaptive staircase procedure.

Test trials relevant for assessing the participants' blur-matching performance at the far test position—the odd-numbered trials—were identical to the trials just described except that the standard image appeared on the far surface and the comparison image appeared on the near surface. Test trials continued until a participant's responses showed at least eight reversals of both the staircase for the near-depth evaluation and the staircase for the far-depth evaluation. Between test trials, participants fixated the central fixation target for 15 s while the surfaces moved to and from the participant under the same conditions as described for Stage 1.

Stage 3 of the experiment, the adaptation training stage, was identical to the baseline training stage with the exception that the images that were mapped onto the surfaces changed over time. Specifically, these images had a blur magnitude that varied as a function of the simulated depth of the surfaces from the participant. Two groups of participants were run. For participants in the "near-blur" group, the images mapped onto the surfaces became more blurry when the surfaces moved nearer to the participant (the images had a blur magnitude of 0 when the surfaces were nearest to a participant and a blur magnitude of 15 when the surfaces were furthest away; this magnitude changed at a constant rate with depth). For participants in the "far blur" group, the images became more blurry when the surfaces moved away from the participant.

Stage 4 (the post-test) was identical to the pre-test stage with the exception that between test trials participants fixated the central fixation target while the surfaces traversed the adaptation range under the same conditions as described for Stage 3.

3. Results

Each participant's point of subjective equality (PSE) is the blur magnitude of the comparison image judged to be more blurry than the standard image with probability equal to 0.5. On both pre-test and post-test, and at both near and far test positions, this value was estimated by averaging the blur magnitudes of the comparison images from the last six reversals of a participant's responses during the adaptive staircase procedure.

Fig. 3 shows the results for a typical participant in the near-blur group (subject ERN). The vertical axis represents the PSE in units of blur magnitude (0 = most blurry image; 15 = least blurry image; 8 = blur of standard image). Light bars represent the participant's PSEs when the standard image was at the far test position (and, thus, the comparison image was at the near test position); dark bars represent the PSEs when the standard image was at the near test position. For the pre-test, the participant's PSEs were close to a value of 8, indicating that the participant matched to the standard image a comparison image of approximately the same blur regardless of whether the standard image appeared at the near or far test position.

For the post-test, however, the participant's responses show a different pattern. When the standard image appeared at the far test position, the participant matched this standard with a comparison image at the near test position that was more blurry than the standard. This pattern of responding can be accounted for as follows. During adaptation training, the participant apparently learned that images in our experimental environment become more blurry as they move near (recall that this participant is in the near-blur group). To compensate for this effect, the participant's visual system adapted by perceptually "de-blurring" near images more than far images. The standard image, which was far, was only mildly de-blurred, whereas the comparison image, which was near, was more significantly de-blurred. Consequently, the participant matched a physically more blurry image to the standard when the standard was at the far test position. When the standard was at the near test position, the opposite pattern emerged: the participant matched the standard with a comparison image that was physically less blurry. Again, this pattern of results is consistent with a model of blur adaptation in which the altered blur–depth relationship experienced during the adaptation training stage changes the perceived blur as a function of depth.

A concise way of summarizing each participant's blur adaptation is to first compare their PSE when the standard appeared at the far test position with their PSE when the standard appeared at the near test position, and then to evaluate whether this far–near difference changed from pre-test to post-test. Fig. 4 shows the differences between each participant's PSE at the far test position (PSE far) and near test position (PSE near) on pre-test (light bar) and post-test (dark bar). Panel A illustrates these PSE differences for the participants in the near-blur group; Panel B shows these values for the participants in the far-blur group. For each group, we

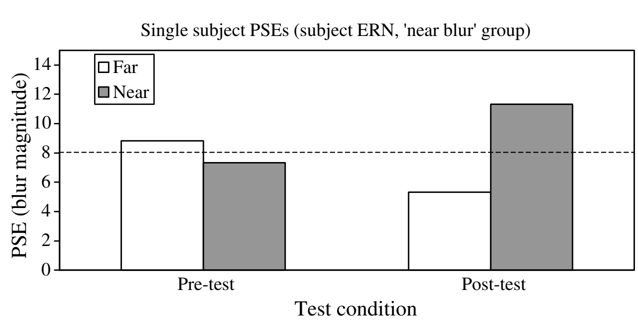

Fig. 3. PSEs of subject ERN on pre-test and post-test when the standard image appeared at the far and near test positions (0 = most blurry; 15 = least blurry; 8 = blur of standard image).

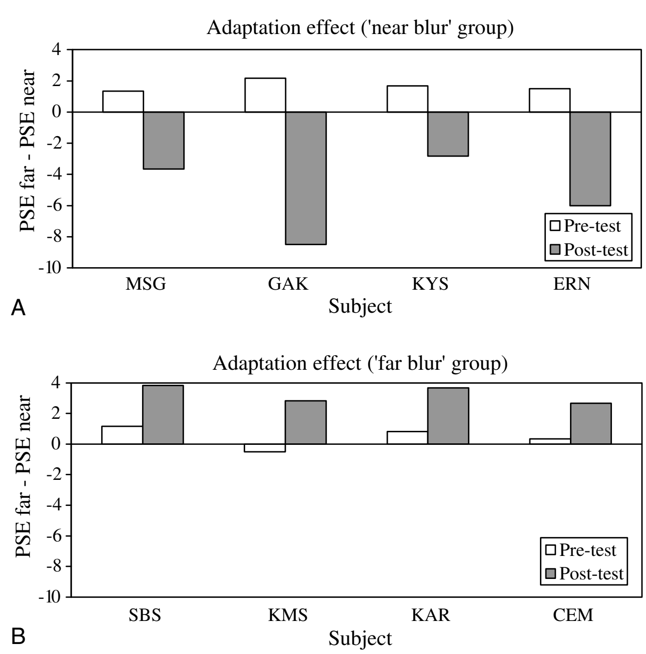

Fig. 4. Differences between each participant's PSEs at the far and near test positions on pre-test and post-test for the four participants in the near-blur group (A) and the four participants in the far-blur group (B).

ran a two-tailed paired *t*-test comparing their pre-test PSE difference (between far and near test positions) with their post-test PSE difference. Both groups showed significant changes from pre-test to post-test ($t(3) = 4.89$, $p < 0.02$ for the near-blur group, and $t(3) = 13.55$, $p < 0.001$ for the far-blur group). Moreover, the directions of the changes are consistent with depth-dependent blur adaptation. We conclude that participants adapted their perception of image blur as a function of surface depth as a result of exposure during the adaptation training stage to images whose blur changed with depth in a novel manner.

4. Discussion

Our results demonstrate that the neural representation of the magnitude of image blur takes into account the depth at which the blur is localized in a scene. We observed that participants can adapt their percept of image blur as a function of the depth at which a blur-matching task is presented. The cause of the adaptation is exposure to an environment with a novel distribution of image blur with respect to depth in a scene. Because the displays in the experiment contained stereo and perspective (e.g., relative size) cues to depth, we cannot draw conclusions about the role of any individual cue in subjects' depth-dependent blur adaptations. It is interesting to speculate, however, that spatial-frequency-tuned stereo channels, such as those studied by Julesz and Miller (1975), may play a part in this phenomenon.

There are at least two hypotheses about the underlying neural mechanisms that could potentially explain these data. The adaptive sensor-gain theory postulates that low-level sensors that process image blur are modulated by a signal conveying information about spatial position, especially depth. This signal dynamically adjusts the gain on the blur sensors so that they adapt their responses based on depth. Thus image blur is represented in 2D coordinates at the sensor level, and this representation is modulated according to visual depth information by an independent signal.

Alternatively, the 3D channel theory postulates that the magnitude of image blur is embedded within the 3D spatial representation. That is, the sensors encode image blur at each local area in the visual scene. There is no need for an independent signal to modulate the gain of the sensors in a depth-dependent manner according to this theory. The present data are not able to differentiate between these two theories of blur perception and adaptation.

The results reported here are compatible with those of several researchers who have recently demonstrated that visual feature detectors can be adapted in a depth-dependent manner (Domini, Blaser, & Cicerone, 2000; He & Nakayama, 1992; Nawrot & Blake, 1989). For example, participants in a study by Blaser and Domini (2002) showed selective adaptation of detectors that were sensitive to both texture and binocular disparity. Similarly, participants in a study by Aslin, Battaglia, and Jacobs (submitted for publication) (Aslin, Jacobs, & Battaglia, 2003) showed contrast adaptation that was sensitive to visual depth. Together, these findings suggest that the visual system is remarkably flexible in adjusting its perception to accurately represent the 3D world.

Acknowledgements

This research was supported by research grants from NSF (SBR-9873477) and NEI (EY-13149) as well as by a NEI Core grant to the Center for Visual Science (EY-01319).

References

Aslin, R. N., Battaglia, P. W., & Jacobs, R. A. (submitted for publication). Depth-dependent contrast gain-control.

Aslin, R. N., Jacobs, R. A., & Battaglia, P. W. (2003). Depth-dependent contrast gain-control. In *Proceedings of the third annual meeting of the vision sciences society, Sarasota, FL*.

Blaser, E., & Domini, F. (2002). The conjunction of feature and depth information. *Vision Research, 42*, 273–279.

Domini, F., Blaser, E., & Cicerone, C. M. (2000). Color-specific depth mechanisms revealed by color-contingent depth after effect. *Vision Research, 40*, 359–364.

He, Z. J., & Nakayama, K. (1992). Surfaces versus features in visual search. *Nature, 359*, 231–233.

Julesz, B., & Miller, J. E. (1975). Independent spatial-frequency-tuned channels in binocular fusion and rivalry. *Perception, 4*, 125–143.

Nawrot, M., & Blake, R. (1989). Neural integration of information specifying structure from stereopsis and motion. *Science, 244*, 716–718.

Webster, M. A., Georgeson, M. A., & Webster, S. M. (2002). Neural adjustments to image blur. *Nature Neuroscience, 5*, 839–840.