Perception, 2002, volume 31, pages 1211 – 1219

DOI:10.1068/p3254

Blur discrimination and its relation to blur-mediated depth perception

George Mather, David R R Smith¶
Experimental Psychology, Biology School, University of Sussex, Brighton BN1 9QG, UK;
e-mail: georgem@biols.susx.ac.uk

Received 2 July 2001, in revised form 14 February 2002;

Abstract. Retinal images of three-dimensional scenes often contain regions that are spatially blurred by different amounts, owing to depth variation in the scene and depth-of-focus limitations in the eye. Variations in blur between regions in the retinal image therefore offer a cue to their relative physical depths. In the first experiment we investigated apparent depth ordering in images containing two regions of random texture separated by a vertical sinusoidal border. The texture was sharp on one side of the border, and blurred on the other side. In some presentations the border itself was also blurred. Results showed that blur variation alone is sufficient to determine the apparent depth ordering. A subsequent series of experiments measured blur-discrimination thresholds with stimuli similar to those used in the depth-ordering experiment. Weber fractions for blur discrimination ranged from 0.28 to 0.56. It is concluded that the utility of blur variation as a depth cue is constrained by the relatively mediocre ability of observers to discriminate different levels of blur. Blur is best viewed as a relatively coarse, qualitative depth cue.

1 Introduction

Human depth perception is supported by a number of cues, including binocular disparity and a range of monocular pictorial cues. Berkeley's *New Theory of Vision*, published in 1709 (see Berkeley 1910) described how very near objects are "seen more confusedly", so that a connection is learnt between the degree of confusion and distance. Berkeley's ideas seem to relate to objects at distances closer than the near point of accommodation. Pentland (1987) identified image blur as a more general source of depth information both in machine vision and in human vision. However, only recently has blur-mediated depth perception become a matter for investigation in psychophysical studies (Marshall et al 1996; Mather 1996, 1997; O'Shea et al 1997; Mather and Smith 2000). Owing to the depth-of-focus limitations of the eye's optics, retinal images of objects nearer or farther than the plane of fixation are blurred by an amount that depends on their relative distance from the fixation plane. Figure 1 plots binocular disparity and estimated retinal image blur as a function of object distance, assuming fixation at 100 cm, based on formulae derived in Mather and Smith (2000).

An obvious difference between the cues is that disparity is signed, but blur is not. This might lead one to assume that it is not possible to determine depth ordering on the basis of blur information alone. However, 'border blur' — the degree of blur at the border between blurred and unblurred regions in the image (Marshall et al 1996; Mather 1996) — offers a possible means of determining depth order using only blur. If the border is blurred, it must be attached to the blurred region, and must therefore be the occluding edge of a nearer (blurred) object. If the border is sharp, it must be attached to the unblurred region, and must therefore be the occluding edge of a nearer (sharp) object. Potentially, then, border blur can disambiguate blur-mediated depth ordering.

¶ Present address: Department of Zoology, University of Cambridge, Downing Street, Cambridge CB2 3EJ, UK

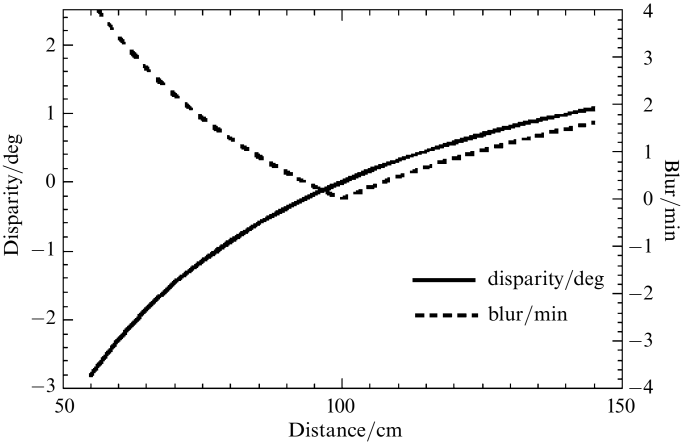

Figure 1. Disparity and blur as a function of distance, assuming fixation at 100 cm, based on formulae provided in Mather and Smith (2000).

Although a number of previous studies have investigated blur discrimination in isolated luminance edges (eg Watt and Morgan 1983), or in regions of texture (eg Mather 1997), there have been no previous attempts to compare blur-mediated depth judgments against blur discriminability in the same stimuli. A comparison of depth and blur judgments may indicate the extent to which the utility of blur-mediated depth cues is limited by the ability of the visual system to discriminate variations in blur. In addition, there have been no previous attempts to measure the discriminability of blur in borders separating regions of texture (second-order blur). In the experiments reported here we sought to (i) measure the effectiveness of region blur and border blur in determining apparent depth ordering; (ii) measure observers' ability to discriminate border blur and region blur in the same stimuli. On this basis, it should be possible to draw inferences regarding the utility of border blur and region blur as cues to depth.

2 Experiment 1: Depth judgments

Observers were shown images containing two textured regions separated by a vertical wavy border. One region was always sharp, the other was blurred (the side containing the blurred region alternated randomly between left and right from presentation to presentation). After each presentation, observers were required to report whether the left-hand or right-hand region appeared farthest away. In 50% of experimental presentations (randomly selected) the border between the two regions was blurred (control trials were also presented in which neither region was blurred but a border was present, and these resulted in no consistent reports of a depth difference between the regions). Example stimuli are shown in figure 2. The degree of blur in the border and in the blurred region of the image were varied independently. Two different viewing distances were employed, to assess whether blur-mediated depth judgments were influenced by fixation distance (note that fixation distance determines the depth interval signalled by a particular degree of retinal blur). If observers make use of border blur when judging depth order, then the blurred stimulus region should appear farther away when the border is sharp, and nearer when the border is blurred.

2.1 *Method*

2.1.1 *Subjects.* Six observers took part in the experiment—both authors and four naïve subjects (three males and three females). All subjects wore appropriate optical corrections, and were aged between 24 and 43 years. Naïve observers were paid for participation.

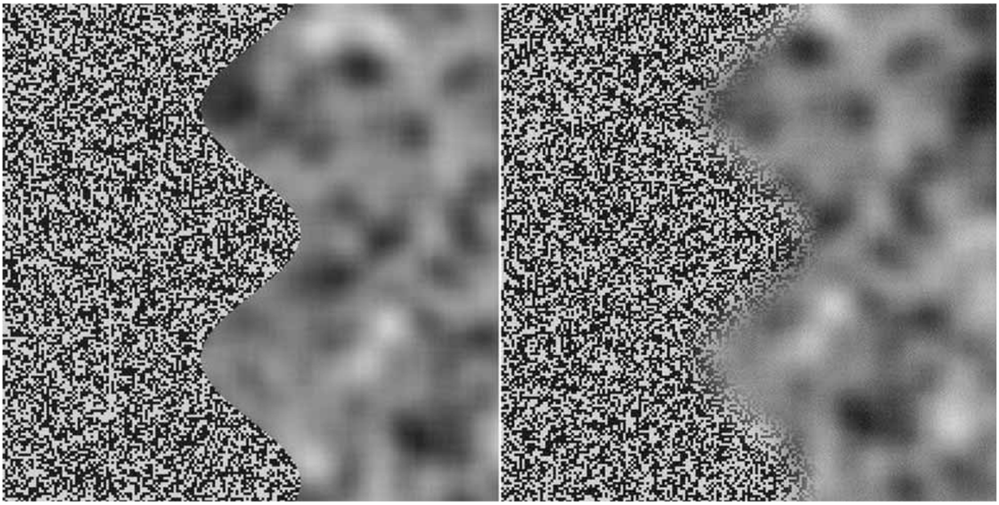

Figure 2. Examples of stimuli used in experiment 1. The left-hand image contains blurred texture on the right, and a sharp border. The blurred region should be seen as far. The right-hand image contains a blurred border, so the blurred region should be seen as near, according to the blur cue.

2.1.2 *Apparatus and stimuli.* Stimuli were generated with a Silicon Graphics O2 workstation and displayed on a GDM-17E21 colour graphic display. The frame rate was 75 Hz, with a horizontal line frequency of 79.8 kHz. Each display pixel measured 0.0226 cm. The minimum and maximum luminance attainable on the monitor was 0.15 cd m^{-2} and 75.0 cd m^{-2} respectively. The power-law nonlinearity of the display was removed by using conventional lookup-table manipulations.

Each stimulus consisted of a 384 pixel × 384 pixel square region of single-pixel binary texture. A vertical sinusoidal border divided the stimulus in half. Two cycles of the sine wave were visible in the stimulus. The sine-wave amplitude was 0.15 as a proportion of stimulus width (see figure 2; all observations were repeated with an amplitude of 0.45, with no difference in results). Four different combinations of region blur, border blur, and viewing distance were employed, as shown in table 1. Blur extent is expressed in terms of the space constant of the Gaussian blur kernel employed. The Michelson contrast of all stimuli was normalised to 0.68 regardless of blurring by using histogram-stretching.

Table 1. Viewing distances, region blur, and border blur used in experiment 1.

Viewing distance/cm	Region blur/min	Border blur/min
57	8	0
57	8	8
57	8	16
57	16	0
57	16	8
57	16	16
114	4	0
114	4	4
114	4	8
114	8	0
114	8	4
114	8	8

The procedure used to create each image can be summarised as follows:

$$I = [S \cdot (G_1 * E)] + [(G_2 * B) \cdot (G_1 * E')] ,$$

where

I is the resultant image,
S is an image containing the texture of one region (to remain sharp),
B is an image containing the texture of the other region (to be blurred),
G_1 and G_2 are Gaussian blur kernels for the edge and region, respectively,
E is a binary image containing the occluding edge as an image mask,
E' is an image containing a bitwise inverted version of E, and
$*$ is the convolution symbol.

This procedure is similar to that employed previously by Marshall et al (1996) but with the added benefit of allowing independent manipulation of the properties of the two image regions (different contrasts, luminances, etc).

2.2 *Procedure*

Subjects completed all observations in two experimental sessions, one at each viewing distance. Each session involved 300 stimulus presentations, comprising 50 randomly ordered presentations of each stimulus in each condition. Each trial involved a single 500 ms exposure of the stimulus, following which the subject pressed one of two response keys to indicate which side of the display (left or right) appeared farther away. No feedback was given as to the correctness or otherwise of the subjects' responses. The display was viewed binocularly without head restraint and with natural pupils. Observers fixated a central fixation mark which was removed from the display during stimulus presentation. The room was kept dark with the only source of illumination coming from the display.

2.3 *Results and discussion*

Figure 3 plots the mean percentage of trials in which the blurred stimulus region was reported as farther away than the sharp region. It is clear that blur in one region of the stimulus resulted in consistent reports that the two regions appeared at different depths, supporting previous reports that image blur acts as a depth cue. However, it is also clear from the figure that neither changing viewing distance nor changing the degree of blur in one region influenced apparent depth ordering.

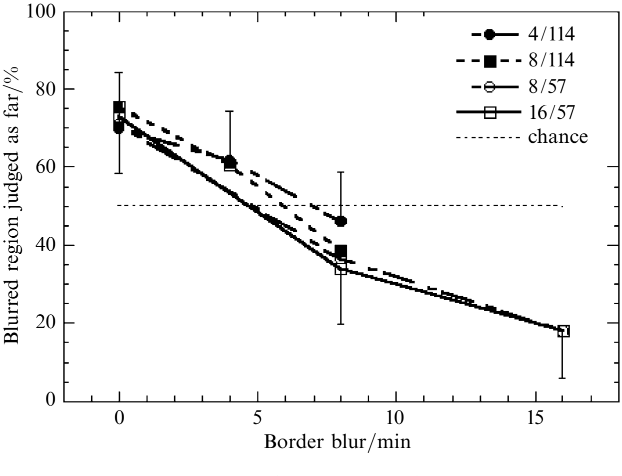

Figure 3. Results of experiment 1, plotting the mean percentage of trials in which the blurred region was judged to be farther than the sharp region as a function of blur at the border between the two regions. Different lines represent different combinations of region blur and viewing distance, eg the 4/114 line represents data from a stimulus with a region blur of 4 min of arc and a viewing distance of 114 cm. Vertical bars represent the SE of the mean.

As expected on the basis of the border-blur cue, the blurred stimulus region was seen as 'far' in all stimulus conditions involving a sharp border (left-most points in figure 3). Reports shifted progressively towards reporting the blurred region as 'near' as the degree of blur in the border increased, as expected if the blurred border is seen as attached to the blurred region. However, when the border was blurred by a moderate amount (around 5 min of arc space constant) apparent depth ordering was ambiguous, as responses fell close to 50%. Why are moderate degrees of border blur ineffective in biasing depth-ordering reports? One possibility is that observers show an inherent bias toward perceiving blurred regions as farther away than sharp regions, and this bias is overruled only by large degrees of border blur. Alternatively, it may be that border blur is simply less detectable than region blur, and therefore not available for use as a depth-ordering cue even at moderate degrees of border blur. Note that in our stimuli the blurred border was between two texture regions, so is in effect a second-order attribute. To address this question we conducted a series of experiments to measure discrimination both of region blur and of border blur using stimuli similar to those used in experiment 1.

3 Experiments 2 – 4: Blur discrimination

3.1 *General method*

3.1.1 *Subjects.* Five subjects took part in the experiments—both authors and three naïve but practised observers.

3.1.2 *Apparatus and stimuli.* The equipment was identical to that used in experiment 1. The viewing distance was 114 cm. Experimental stimuli are illustrated in figure 4. *Experiment 2* measured discrimination of blur in luminance borders. The stimulus (left-hand panel in figure 4) contained a wavy luminance border separating regions of low luminance and high luminance (11.8 cd m^{-2} and 61.5 cd m^{-2} respectively; Michelson contrast 0.68). *Experiment 3* measured discrimination of blur in texture borders. The texture on one side of the border consisted of vertically elongated random binary noise, and the texture on the other side of the border consisted of horizontally elongated random binary noise (aspect ratio of 1 : 50; middle panel of figure 4). The two textures were matched for luminance and contrast with the luminance-border stimulus used in experiment 1. The textures were always rendered as sharp and unblurred, but the border between them was blurred under the procedure summarised earlier. *Experiment 4* measured discrimination of region blur. The same anisotropic textures were used as in experiment 3, except that now the border was always rendered as sharp, and the texture in each region was blurred (right-hand panel of figure 4).

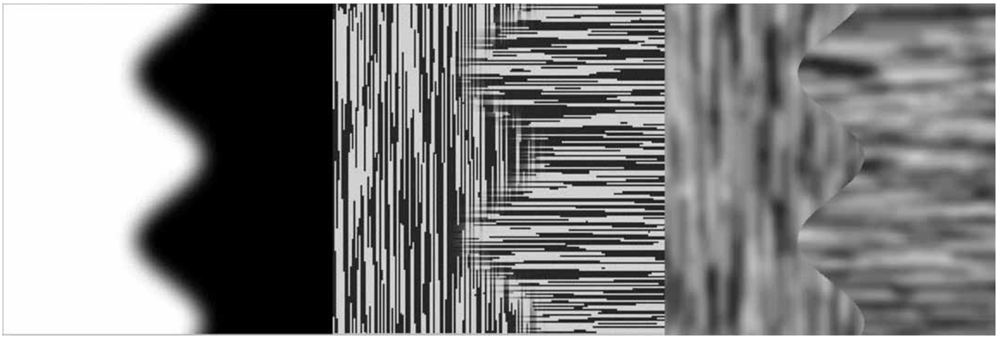

Figure 4. Stimuli used to measure blur discrimination in experiments 2 to 4. Left—luminance-border blur; middle—texture-border blur; right—region blur.

Each stimulus was rendered with five different degrees of reference blur, having space constants of 0, 1, 2, 4, and 8 min of arc. The stimuli were presented foveally in a square patch (4.36 deg along each side) with no spatial and temporal windowing. Each experiment measured the observer's ability to discriminate small increments in blur about each of these reference blurs.

3.1.3 *Procedure.* A self-paced two-temporal-interval forced-choice (2AFC) paradigm was used in conjunction with the method of constant stimuli to obtain full psychometric functions (at least 100 observations for each psychometric curve measured at five or more different points). Each trial was initiated by pressing a mouse button, and consisted of two temporal intervals each lasting 500 ms. The two temporal intervals were separated by a 1 s period, during which the display was reset to a uniform mean luminance of 37.5 cd m^{-2} and the central fixation mark was redisplayed. The stimulus in one temporal interval contained a specific reference border blur, and the stimulus in the other temporal interval contained an increment or decrement from this blur. Stimuli were presented against a uniform grey background at mean luminance (37.5 cd m^{-2}). The background filled the entire monitor. The observer's task was to signal, by pressing the appropriate mouse button, which temporal interval contained the most-blurred stimulus. Observers were not given feedback as to the correctness or otherwise of their responses.

Data for each experiment were gathered over multiple sessions. Within a session, reference blur was selected at random from trial to trial, with the constraint that no reference would be presented for the $(n+1)$th time until all references had been presented n times. Blur increment was selected at random from a set of five values associated with each reference blur value (chosen on the basis of earlier experiments to provide a well-sampled psychometric function), with the constraint that no increment would be presented for the $(n+1)$th time until all comparisons had been presented n times. Each blur increment was displayed four times (twice in the first interval, once on the left and once on the right, and twice in the second interval). The temporal interval and side of presentation of the reference image was selected at random. Each observer progressed through the sequence of experiments in a different order. Data were accumulated for each subject over thirty experimental sessions until twenty trials had been presented for each stimulus in each experiment.

3.2 *Results*

Cumulative Gaussians were fitted to the data of each observer using Probit analysis (Foster and Bischof 1997). The reciprocal of the slope of the fitted function at the 50% point was used to estimate the blur-discrimination performance of each observer for each reference blur.(1) Average blur discrimination for the five observers in each of the experiments is shown in figure 5. Blur-discrimination performance follows a typical 'dipper' function (eg Mather 1997) where performance initially improves for small reference blurs and then rises in proportion to reference blur. The fitted function in figure 5 assumes that blur increment threshold (Δb) is determined by the total blur present in the stimulus (σ_t) and the Weber fraction (w) for blur discrimination:

$$\Delta b = w\sigma_t \,. \qquad (1)$$

Total blur contains contributions from extrinsic blur (σ_e) and intrinsic blur (σ_i). Extrinsic blur corresponds to the reference blur space constant. Intrinsic blur represents blur contributed by the visual system, being a combination of optical blur in the

---

(1) For the sharp (0 blur) reference the data for comparison blurs greater than 0 were reflected back for the purposes of fitting a curve to a full psychometric function. This allowed us to find the reciprocal of the slope at the 50% point.

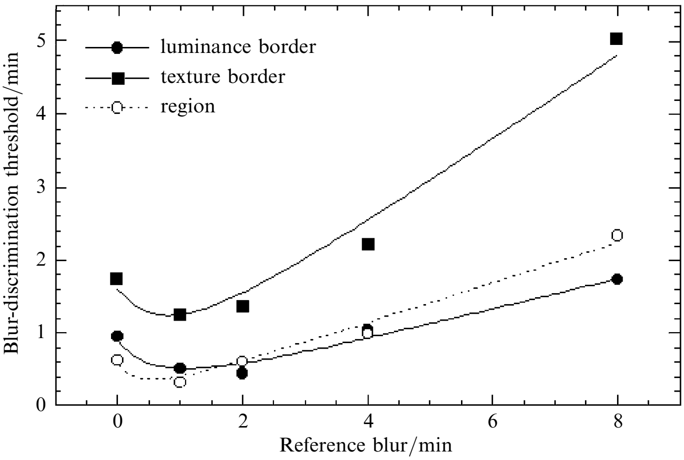

Figure 5. Results of experiments 2 – 4, plotting mean blur-discrimination threshold as a function of reference blur for the three stimuli illustrated in figure 4. Each datum point is calculated from at least 500 trials over five observers. Smooth curves assume a constant Weber fraction for blur discrimination, allowing for the presence of a constant intrinsic neural blur. The minimum of the function estimates the space constant of the intrinsic blur. Standard errors on individual data points have been omitted for clarity. Average SEs in each condition were: luminance border—± 0.17; texture border—0.23; region—0.20.

eye and neural blur in receptive fields. According to the addition of variances rule, total blur is given by:

$$\sigma_{\mathrm{t}} = (\sigma_{\mathrm{i}}^2 + \sigma_{\mathrm{e}}^2)^{-1/2}\,. \tag{2}$$

The 'dipper' shape arises because, at near-zero reference blurs, relatively large extrinsic blur space constants are required to overcome the contribution of the intrinsic blur space constant. As reference blur increases beyond the value of the intrinsic space constant, the contribution of the latter diminishes. The minimum of the dipper function corresponds to the value of the intrinsic blur space constant (Watt 1988; Mather 1997). Table 2 shows the value of the Weber fraction and the intrinsic blur space constant for the best-fitting function in each experiment, along with the coefficient of determination.

Table 2. Weber fractions and intrinsic blur space constants for the best-fitting function in each experiment.

	Weber fraction	Intrinsic blur/min	r^2
Luminance border	0.213	1.310	0.96
Texture border	0.590	1.297	0.97
Region	0.278	0.729	0.98

3.3 *Discussion*

Equation (1) provides a convincing account of the data, as indicated by the high coefficients of determination in table 2. Several important points emerge from a consideration of the other values in table 2. First, the estimated internal blur space constant for region blur discrimination of 0.729 min of arc agrees well with values in the literature, such as the estimate of 0.71 min of arc reported for blur discrimination of fractal textures (Mather 1997). Second, the intrinsic blur space constant for border blur discrimination, both luminance and texture defined, is appreciably larger at around 1.3 min of arc. This indicates that blur discrimination in these larger-scale structures is mediated by receptive fields with large space constants. Third, Weber fractions for texture-border blur discrimination are much larger than those for luminance-border discrimination, indicating that subjects found it very difficult to discriminate different degrees of blur in texture borders. It could be argued that performance with texture borders should be compared against that with low contrast luminance borders rather than the high contrast borders used so far. We therefore conducted a supplementary experiment to measure luminance-border blur discrimination at a range of contrasts between 0.1 and 0.8. Three subjects took part in the experiment—one author and two naïve observers. All details were identical to those for the previous experiments, except that only the luminance-border stimulus was used, and five different contrasts were presented in different experimental sessions (0.1, 0.2, 0.4, 0.6, and 0.8).

Figure 6 shows the mean blur-discrimination threshold as a function of reference blur, with contrast as the parameter. An eight-fold variation in contrast had very little effect on discrimination performance. The Weber fraction at a contrast of 0.1 was 0.233, and the Weber fraction at a contrast of 0.8 was 0.221. We conclude that the poor discriminability of blur in texture borders cannot be likened to the effect of reducing the contrast in intensity borders. Poor performance with texture borders must represent an inherent property of the mechanisms mediating their detection, such as that proposed by Bergen and Landy (1991).

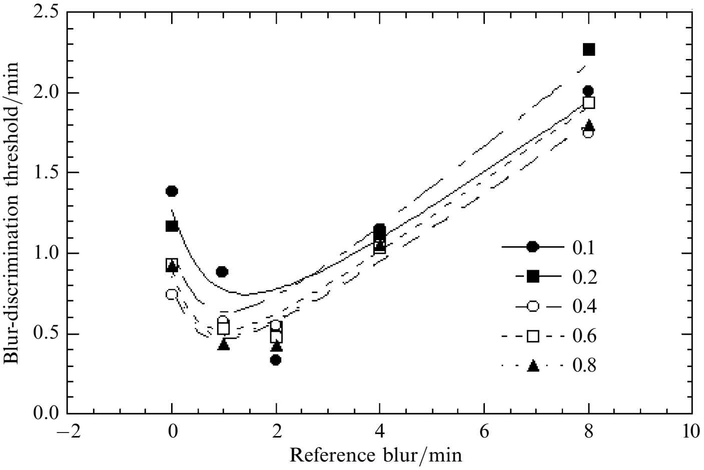

Figure 6. Results of experiment 5, plotting mean blur-discrimination threshold as a function of reference blur, for five different levels of contrast. Each datum point is calculated from at least 300 trials over three observers. Standard errors on individual data points have been omitted for clarity. Average SE was 0.173.

4 General discussion

The results of these experiments indicate that image blur is best viewed as a qualitative cue to pictorial depth. In experiment 1 we found that moderate degrees of texture-border blur were ineffective in biasing depth-ordering reports. Border blur influenced depth ordering only at extreme values (ie either no blur at all or a large degree of border blur). Two possible explanations were offered: (i) subjects have an inherent bias toward perceiving blurred regions as farther away than sharp regions; or (ii) border blur is less detectable than region blur, and therefore not available as an ordering cue at moderate degrees of blur. Subsequent experiments on blur discrimination favoured the second explanation: texture-border blur is much less detectable than region blur (see table 2).

More generally, it seems that the effectiveness of image blur as a depth cue is limited by the ability of the visual system to discriminate small changes in blur. Weber fractions for blur discrimination ranged from about 0.21 to 0.56 in our experiments. The smallest Weber fractions were obtained for discrimination of luminance-border (or first-order) blur, while the largest fractions were obtained for discrimination of texture-border (or second-order) blur. For comparison, Weber fractions for stereopsis are typically below 0.1 (Badcock and Shor 1985).

The poor discriminability of blur indicated by these and other experiments contrasts sharply with the high sensitivity of the accommodative system. Kotulak and Schor (1986) found that an accommodative response can be elicited by a blur stimulus that is below the threshold for blur perception. They concluded that the accommodative system uses a mechanism that does not rely on perceptible levels of blur.

Blur and binocular disparity offer closely coupled quantitative cues for depth perception (Mather and Smith 2000). However, small variations in blur extent are much less visible than small variations in disparity. Consequently, blur should be viewed as a relatively coarse qualitative cue to depth, akin to traditional pictorial cues, rather than a precise quantitative cue akin to disparity.

Acknowledgments. This research was supported by a grant from the Engineering and Physical Sciences Research Council, UK.

References

Badcock D, Shor C, 1985 "Depth-increment detection function for individual spatial channels" *Journal of the Optical Society of America A* **2** 1211 – 1215

Bergen J R, Landy M S, 1991 "Computational modelling of visual texture segregation", in *Computational Models of Visual Processing* Eds M S Landy, J A Movshon (Cambridge, MA: MIT Press) pp 253 – 271

Berkeley G, 1910 *A New Theory of Vision and Other Writings* (London: Dent)

Foster D H, Bischof W F, 1997 "Bootstrap estimates of the statistical accuracy of thresholds obtained from psychometric functions" *Spatial Vision* **11** 135 – 139

Kotulak J, Schor C, 1986 "The accommodative response to subthreshold blur and to perceptual fading during the Troxler phenomenon" *Perception* **15** 7 – 15

Marshall J, Burbeck C, Ariely D, Rolland J, Martin K, 1996 "Occlusion edge blur: A cue to relative visual depth" *Journal of the Optical Society of America A* **13** 681 – 688

Mather G, 1996 "Image blur as a pictorial depth cue" *Proceedings of the Royal Society of London, Series B* **263** 169 – 172

Mather G, 1997 "The use of image blur as a depth cue" *Perception* **26** 1147 – 1158

Mather G, Smith D R R, 2000 "Depth cue integration: Stereopsis and image blur" *Vision Research* **40** 3501 – 3506

O'Shea R, Govan D, Sekuler R, 1997 "Blur and contrast as pictorial depth cues" *Perception* **26** 599 – 612

Pentland A P, 1987 "A new sense for depth of field" *IEEE Transactions on Pattern Analysis and Machine Intelligence* **9** 523 – 531

Watt R J, 1988 *Visual Processing* (Hove, Sussex: Lawrence Erlbaum Associates)

Watt R J, Morgan M, 1983 "The recognition and representation of edge blur: Evidence for spatial primitives in human vision" *Vision Research* **23** 1465 – 1477