
The use of image blur as a depth cue

George Mather

Experimental Psychology, School of Biology, University of Sussex, Brighton BN1 9QG, UK

Received 18 October 1996, accepted 16 April 1997

Abstract. Images of three-dimensional scenes inevitably contain regions that are spatially blurred by differing amounts, owing to depth-of-focus limitations in the imaging apparatus. Recent perceptual data indicate that this blur variation acts as an effective cue to depth: if one image region contains sharply focused texture, and another contains blurred texture, then the two regions may be perceived at different depths, even in the absence of other depth cues. Calculations based on the optical properties of the human eye have shown that variation in blur as a function of depth follows the same course as variation in binocular disparity with depth. Computational modelling has shown that the effect of blur on single-step edges is very similar to its effect on random fractal patterns, because the two stimuli have similar Fourier amplitude spectra. Blur discrimination thresholds for the two stimuli were also very similar, and could be predicted by a model based on high-spatial-frequency discrimination. A comparison of blur discrimination thresholds with the range of binocular stereopsis indicates that blur and disparity cues cover different distance ranges: stereopsis is most effective for distances relatively close to fixation, while blur information should be more effective for larger distances.

1 Introduction

Human depth perception is mediated by a number of different cues. Qualitative information about depth ordering of objects in the image is provided by such cues as interposition, relative size, and height in the visual field. Quantitative information about the magnitude of depth differences between image features is provided primarily by binocular stereopsis and motion parallax. Convergence and accommodation offer nonvisual metrics for absolute distance [see Howard and Rogers (1995) for a detailed review of depth cues]. Recent perceptual research demonstrates that the list of known depth cues should also include image blur. Images of three-dimensional scenes inevitably contain variations in spatial blur, due to depth-of-focus limitations in the imaging apparatus. When an observer is viewing a three-dimensional scene containing objects at different distances, image detail (eg texture markings) arising from objects at accommodation distance will be in sharp focus (assuming no refractive error). Image detail arising from objects at other distances will be blurred owing to depth-of-focus limitations, to an extent that depends on their distance from fixation. The degree of blur at the border between less and more blurred regions can be used to specify whether the blurred region is nearer or farther than fixation distance. If the sharply imaged object is nearer and therefore occluding the blurred object, then the border itself should be sharp since it is attached to the sharp object. If the sharply imaged object is farther, and therefore being occluded by the blurred object, then the border should be blurred. The two simple images in figure 1 demonstrate some of these effects.

Psychophysical evidence for the use of image blur in depth perception has recently been reported by Marshall et al (1996) and by Mather (1996). Both papers described experiments on ambiguous figure-ground stimuli, containing two regions of texture separated by a wavy boundary. When texture in one region was blurred while texture in the other region was sharp, the two regions appeared to be at different depths. When the boundary itself was sharp, the sharp texture was seen as nearer, and occluding the blurred texture region, but when the boundary was blurred, the blurred texture was seen as nearer.

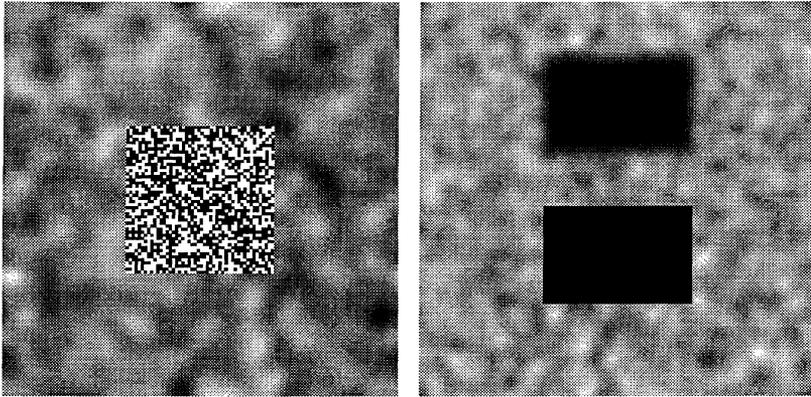


Figure 1. Blur as a depth cue in random patterns. In the left-hand image a region of sharp binary texture appears to stand forward of the blurred binary texture surrounding it. In the right-hand image, the background again contains blurred binary texture. The lower black rectangle appears nearer than the upper rectangle, because it has sharply defined edges.

The main aims of this paper are to consider: (i) the physical properties of image blur as a cue for human depth perception; and (ii) how estimates of image blur are made available for depth perception.

2 Information available in image blur

The following general expression relates the distance d of a point from a lens to the radius s of its blurred image (Pentland 1987):

$$d = \frac{Fv}{rv - F(r + s)} \quad (1)$$

where F is focal length, r is lens aperture radius, and v is the distance of the image plane from the lens. Note that this expression does not distinguish between points that are farther and nearer than the point of focus. In principle, if the values of F , r , and v are known and a measure of image blur s is available, then absolute distance can be calculated (provided that the nearer/farther ambiguity is resolved).

Equation (1) can be used to predict retinal blur as a function of distance, on assuming typical values for the optical parameters of the human eye ($r = 1.5$ mm, $v = 16$ mm). Figure 2a (solid line) shows predicted blur as a function of distance when the lens is accommodated on a point 1 m away (so $F = 63.5$ D), and figure 2b shows predictions for accommodation at 4 m ($F = 62.75$ D).

It should be noted that equation (1) and figure 2 assume geometrical optics. Owing to diffraction and aberrations, blur in retinal images never actually declines to zero, but reaches a minimum of approximately 0.7 min arc (eg Campbell and Gubisch 1966). For comparison, figure 2 also shows the magnitude of binocular disparity over the same distances (broken lines; calculated as the difference in vergence angle between fixation and the specified distance, with inter-pupillary distance taken as 63 mm). Blur increases smoothly as distance from fixation increases, and follows the same curve as binocular disparity. The absolute values of the two cues are very different (note the y axes in figure 2), but the significance of this difference depends on the sensitivity of the visual processes encoding the cues, as discussed later. Apart from focal length and blur extent, the other free parameter in equation (1) is aperture radius. Pupil diameter in humans can range between approximately 2 mm and 8 mm, and depends on a number of influences, both visual and nonvisual (see Thompson 1981).

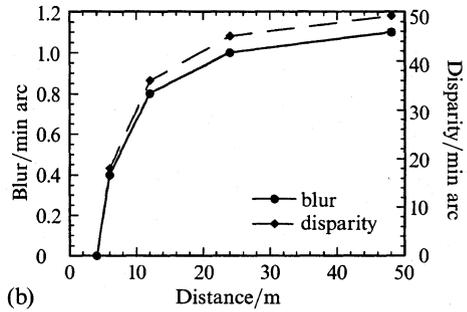
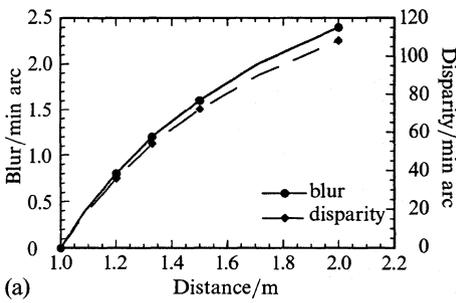


Figure 2. Retinal image blur (solid lines) and binocular disparity (broken lines) as a function of distance. Image blur was calculated with the use of equation (1); see text for details. Calculations for fixation at (a) 1 m and (b) 4 m.

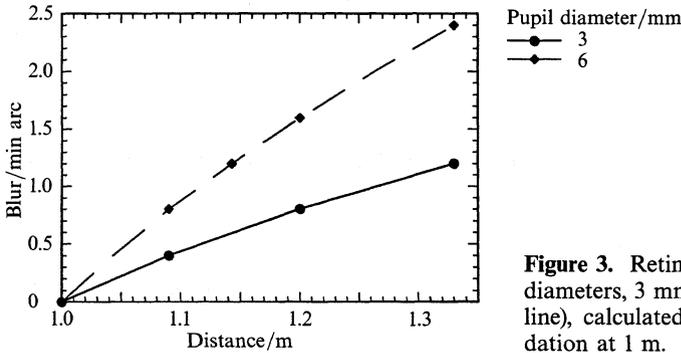


Figure 3. Retinal image blur for two pupil diameters, 3 mm (solid line) and 6 mm (broken line), calculated as in figure 2, for accommodation at 1 m.

Figure 3 illustrates how pupil diameter affects predicted blur extent. The degree of blur at any one distance is proportional to pupil diameter.

So far, we have considered optical blur as it applies to a single point in the image. How does blur affect the appearance of more complex visual stimuli? In considering the effect of blur, a Gaussian point spread function will be assumed:

$$G(x, y) = \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \tag{2}$$

The space constant of the Gaussian σ specifies blur extent [Pentland (1987) showed that Gaussian blur offers a close approximation to blur induced by depth-of-field limitations].

Previous computational and psychophysical research has concentrated on the local effect of blur on individual intensity edges (eg Watt and Morgan 1983; Pentland 1987; Subbarao and Gurumoorthy 1988). Figure 4a illustrates the effect of applying Gaussian blur to a step intensity edge.

As blur increases, the rate of change of intensity across the edge declines, and the width of the graded intensity region expands. Figure 4b shows Fourier amplitude spectra for the edges in figure 4a. For a sharp edge, amplitude declines as a function of $1/f$ so, when plotted on log-log axes, amplitude falls off linearly with a slope of -1.0 . Gaussian blur selectively removes high-frequency components from the spectrum.

On a more global scale, recent research has highlighted the consistent statistical properties of natural images (eg Field 1987; Knill et al 1990; Tadmor and Tolhurst 1994; van der Schaaf and van Hateren 1996), which have scale-invariant fractal characteristics. In the Fourier domain, such images (containing, for example, vegetation, water, rocks, mountains) have amplitude spectra that, on average, fall off roughly as a function of $1/f$, though there is a degree of variability. Just as in the case of single-step edges, amplitude falls off linearly with a slope of -1.0 when plotted on log-log axes.

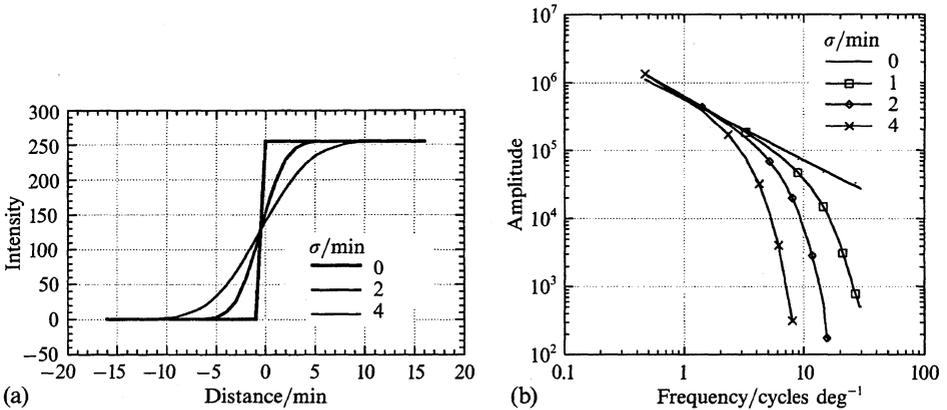


Figure 4. (a) Profiles for a step intensity edge that is unfiltered ($\sigma = 0$), blurred with a Gaussian space constant of 2 min, or blurred with a Gaussian space constant of 4 min. (b) The Fourier amplitude spectrum of the edge for different Gaussian blur space constants.

As a way of modelling the effect of blur on $1/f$ images, Gaussian spatial blur was applied to random fractal texture patterns that have the required $1/f$ amplitude spectra (Brownian fractals; see Peitgen and Saupe 1988). Figure 5 illustrates such a pattern (5a, top left) and its amplitude spectrum on log–log coordinates (5b, heavy line).

Figure 5 also shows similar patterns and amplitude spectra after application of Gaussian blurring functions with different space constants.

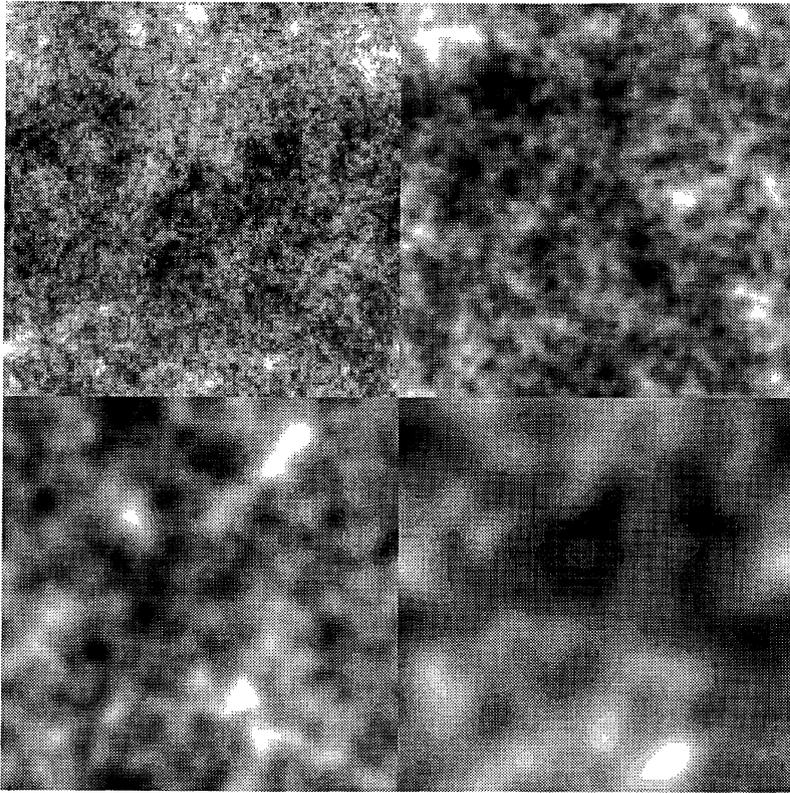
There are close similarities between the effect of blur on single edges and the effect of blur on fractal texture, because in both images Fourier amplitude falls off as a function of $1/f$. Blur removes energy at high frequencies, and the relationship between the highest spatial frequency remaining in the pattern and the space constant of the blurring filter conforms to a power law with an exponent of -1.0 . Note that variations in blur extent amount to variations in the spatial scale of the image or, equivalently, to translations along the frequency axis of the Fourier amplitude spectrum. The edge profiles in figure 4 differ in terms of x -axis dilation, and the textures in figure 5 differ in terms of spatial magnification.

3 Analysis of image blur by the human visual system

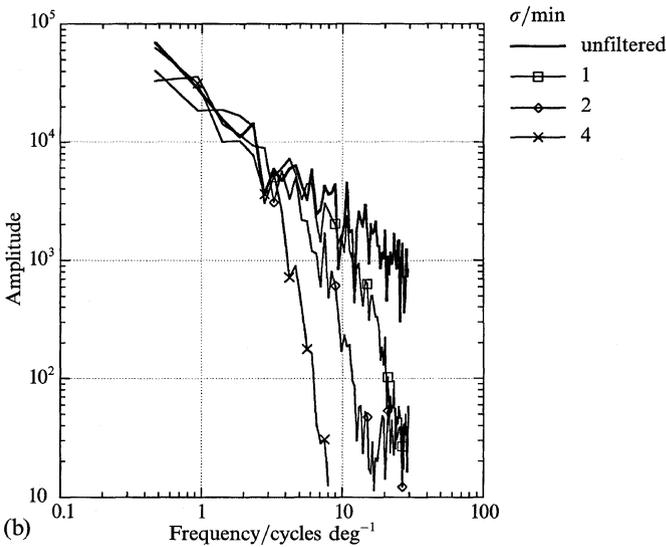
It is clear that blur plays a role in depth perception (Marshall et al 1996; Mather 1996), and there is a sound physical basis for the cue [equation (1) and figure 1], but only a small number of papers have been concerned with blur estimation by the human visual system. Several theoretical schemes have been proposed, both in the computational-vision literature and in the psychophysical literature, but so far only schemes in the former have been developed in the context of depth estimation.

3.1 Multiple images

This method involves a comparison of two images that are identical except for having been formed with different lens apertures or focal lengths. The resulting point-by-point differences between the images can be used to compute blur. Pentland (1987) and Ens and Lawrence (1993) describe implementations of this method that rely on a comparison between one image formed with a small aperture (a pinhole in Pentland's method) and another formed with a larger aperture. Pentland speculated on two ways in which the human visual system might implement this method. First, "the depth of field for the red–green retinal cells is different from that for the blue retinal cells, because of one diopter of chromatic aberration in the lens". It seems extremely unlikely that such information is available to human vision, since there are so few blue photoreceptors available (eg Marc and Sperling 1977), probably because of problems caused by chromatic



(a)



(b)

Figure 5. (a) Blurred random fractal patterns (each 128×128 pixels). Top left, unblurred; top right, Gaussian space constant equal to 1 pixel; bottom left, space constant equal to 2 pixels; bottom right, space constant equal to 4 pixels. (b) Fourier amplitude spectra for random fractal patterns similar to those in (a), for different Gaussian blur space constants. These 1-D spectra represent information in horizontal slices across the 2-D patterns.

aberration, and photoreceptor outputs are pooled at the ganglion-cell level. Pentland further suggested that human implementation of this method relies on the observation that “the focal length of the human eye is constantly varying in a sinusoidal fashion at a frequency of about 2 Hz” (as first reported by Campbell et al 1959). This oscillation would—he argues—result in depth-of-focus oscillation sufficient to estimate distance. However, the magnitude of the oscillation appears to be only at or below the threshold for blur detection, perhaps because it plays a role in maintaining accommodation (eg Kotulak and Schor 1986; Winn et al 1989), so it seems unlikely that the oscillation would generate sufficient blur for depth perception. Of course, larger voluntary changes in accommodation do occur, and result in point-by-point differences in blur extent that depend on distance, and these could play a role in depth estimation. However, blur can be an effective depth cue even in single views of photographic or pictorial images, such as those in figure 1, so blur variation induced by focus changes is certainly not necessary for blur-mediated depth.

3.2 *Edge-based schemes*

A number of proposed methods estimate blur in a single intensity edge using space-domain filtering. As shown in figure 4, the width of the intensity gradient across the edge widens as blur extent increases. Edge-based blur estimation schemes almost universally begin with the application of spatial-differentiating filters to the edge profile. The output of a first-derivative filter will reach a peak at the location of the edge (point of inflexion in the blurred edge profile). The output of a second derivative filter will cross zero at the edge location, and will have a peak and a trough on the dark and light sides of the edge, respectively. The gradient of the second-derivative at the zero-crossing will decrease, and the distance between positive and negative peaks on either side will increase, as blur extent increases. A simple measure of blur extent would be the maximum rate of change of intensity across the edge (eg slope of the zero-crossing in the second derivative). Schemes in the computational literature attempt to measure zero-crossing slope in various ways (eg Pentland 1987; Subbarao and Gurumoorthy 1988). A simple measure of zero-crossing slope would confound edge contrast with blur extent, and human observers are immune to this problem when estimating blur width (Watt and Morgan 1983; Georgeson 1994). Watt and Morgan proposed that the visual system estimates blur by measuring the distance between adjacent peaks and troughs in the output of the second-derivative filter. More recently, Georgeson proposed that blur is computed by taking the square root of the ratio of the first to third derivatives of the blurred edge at the zero-crossing of the second. The response of centre-surround receptive fields in the visual system approximates the second derivative of the image, so this aspect of both models is plausible. Georgeson’s model requires receptive fields that compute first and third derivatives as well.

Watt and Morgan (1983) measured blur discrimination thresholds for isolated intensity edges blurred with the use of three different functions (Gaussian, rectangular, and cosinusoidal blur). Thresholds varied between about 0.2 min arc and 2 min arc, and Watt and Morgan argued that results agreed with predictions based on the peak-trough distance in the second derivative of the stimulus. Georgeson (1994) reported an experiment in which subjects matched the apparent blur of a Gaussian-blurred edge against the perceived blur of sine gratings (both extended patches and truncated gratings). Matched edge blur was directly proportional to grating frequency. Predictions based on the ratio of the first to third derivatives, and on peak-trough separation, could both account for the proportional relationship, but the fit of the former was closer than that of the latter.

3.3 *Frequency-based schemes*

This approach to blur estimation rests on the observation that image blur removes high spatial frequencies from the Fourier amplitude spectrum. It is well known that receptive fields at a range of sizes (ie preferred spatial frequencies) exist at each location in the visual field. De Valois and De Valois (1988) suggested that blur at different locations could be determined by comparing "the relative amount of power at high frequencies present in various sub-regions" (page 293). From the frequency spectra in figures 4 and 5 it is clear that both single edges and natural images as a whole have $1/f$ spectra. According to Field (1987), an array of spatial-frequency-tuned receptive fields that had equal bandwidth on an octave scale would be equally activated by $1/f$ spectra, since such spectra contain equal energy in equal octave intervals. The actual bandwidth of spatial-frequency-tuned channels in the visual system is fairly constant on an octave scale, though there is evidence for a gradual fall in bandwidth as centre frequency rises (De Valois et al 1982; Georgeson and Harris 1984). Consequently, in the absence of blur there may be a gradual decline in channel response as preferred frequency increases (though recall that actual spectral slope does vary from image to image; Tadmor and Tolhurst 1994; van der Schaaf and van Hateren 1996). However, blur produces a sharp cut in energy at a frequency that is inversely proportional to blur extent. Visual channels below the cut-off frequency would be unaffected, but there should be a rapid decline in activation at higher frequencies. Blur estimation may therefore be based on a comparison of the outputs of spatial-frequency-tuned channels in the visual system.

The existence of the size aftereffect (Blakemore et al 1970) and simultaneous spatial-frequency-shift effect (Klein et al 1974) is consistent with the idea that apparent spatial scale depends on the relative activity of frequency-tuned channels. In addition, Brown and Weisstein (1988) reported a link between the apparent size of sine gratings and depth perception. Using stimuli in which alternating horizontally layered slices were filled with vertical sine gratings at different frequencies, they found that slices containing the higher-frequency grating appeared closer in depth than those with the lower-frequency grating. As noted by Marshall et al (1996), the apparent nearness of the high-frequency grating may result from using slices with sharp edges.

Frequency-based estimation of blur in a region of the retinal image may be based on the same inhibitory interactions between frequency-tuned channels that have been invoked to explain shifts in perceived spatial frequency. It was shown earlier that the highest spatial frequency remaining in the image is inversely proportional to blur extent. Mutual inhibition between channels covering a particular region of the image would result in a relatively high level of activity in the channel responding to the highest frequency present, because that channel would receive no inhibition from higher-frequency channels. Smallman et al (1996) recently presented evidence for channels tuned to frequencies higher than $34 \text{ cycles deg}^{-1}$ that may mediate discrimination of relatively small blur extents (figures 4 and 5 show that blurring with a space constant of 1 min arc effectively leaves no energy at frequencies over $30 \text{ cycles deg}^{-1}$).

According to the frequency-based account, blur discrimination should be closely linked to spatial-frequency discrimination. Campbell et al (1970) found that spatial-frequency difference thresholds obtained with sine gratings are a fixed proportion of the criterion frequency (ie a power law with an exponent of 1.0). Watt and Morgan (1983) reported a power-law exponent of 1.5 for edge blur discrimination, and on this basis rejected a frequency-based account. However, Paakkonen and Morgan's (1994) data on blur discrimination were consistent with an exponent of 1.0. Paakkonen and Morgan attribute Watt and Morgan's steeper function to the use of short edge segments, and their observations confirmed that the power-law exponent is affected by edge length.

Tadmor and Tolhurst (1994; see also Tolhurst and Tadmor 1997) performed an experiment in which observers were required to discriminate the spectral slope of natural

and random fractal images. Images were filtered in the Fourier domain to have spectral slopes ranging from -0.4 to -1.6 ; $1/f$ patterns have a spectral slope of -1.0 , as noted earlier. Observers were shown trios of images, one of which differed from the other two in its spectral slope, and were required to select the 'odd one out'. Slope discrimination values were in the range 0.05 – 0.15 . Tadmor and Tolhurst reported that, subjectively, natural images appear blurred at high spectral-slope values and, in a physical sense, "the task of discriminating changes in α [slope] is analogous to the discrimination of changes in the degree of blur" (1994, page 552). Both blur and slope manipulations attenuate the amplitude of high-frequency components in the spectrum.

Although psychophysical data are available on blur discrimination in single intensity edges, no data are available for more complex stimuli. Therefore a blur discrimination experiment was conducted with random fractal patterns, to establish whether data are comparable with those obtained with single edges.

4 Experiment: blur discrimination in random fractal patterns

4.1 Methods

4.1.1 *Subjects*. Four naive observers took part in the experiment.

4.1.2 *Apparatus*. Stimuli were generated on a PC-compatible computer equipped with a high-resolution graphics card, and displayed on a Hitachi 14MVX colour monitor. Responses were recorded by means of buttons on a standard PC mouse.

4.1.3 *Stimuli*. A large library of random fractal images was compiled for the experiment. Each image was generated as follows. First, a 128×128 pixels 'white noise' image was generated (the grey level of each pixel was randomly selected with the use of a Gaussian random number generator; see Peitgen and Saupe 1988). The 2-D discrete Fourier transform of the image was computed by conventional methods (the amplitude spectrum of such an image is flat ie with equal power at all frequencies). The amplitude of each frequency component in the spectrum was adjusted to generate $1/f$ noise (ie amplitude was scaled to fall off as a function of $1/f$). If necessary, the amplitude spectrum was also filtered by using a Gaussian low-pass function before the inverse Fourier transform was computed to generate a blurred random fractal pattern. All images were normalised to the same mean intensity (25.4 cd m^{-2} on the display monitor) and rms contrast (0.2), and therefore had the same total spectral power. Examples of resulting images are shown in figure 5. The power-law nonlinearity of the display was compensated by means of conventional look-up table manipulation. Images subtended $2.14 \text{ deg} \times 2.14 \text{ deg}$ square at the 268 cm viewing distance.

4.1.4 *Procedure*. Stimuli were presented as pairs of images side-by-side (0.54 deg between inner edges) against a mean-luminance background for a fixed presentation time of 0.5 s , with a small fixation cross drawn midway between them. After each presentation, the subject identified the image that appeared more blurred by pressing a mouse button. One member of each image pair (the reference image) had a Gaussian blur space constant of either 0 (unblurred), 1 , 2 , or 4 min . The degree of blur present in the other member of each pair (the comparison image, which was always a different random pattern from the reference image) varied from trial to trial. The left/right location of the two images also varied randomly. Different combinations of blur were shown in random order over two sessions according to the method of constant stimuli, until each of the four observers had accumulated 40 responses to each combination.

4.2 Results and discussion

Psychometric functions are plotted in figure 6a. Thresholds, defined as the slope of the function at its midpoint (calculated with the use of logistic regression; Berkson 1953), are plotted in figure 6b. For comparison, the figure also shows discrimination

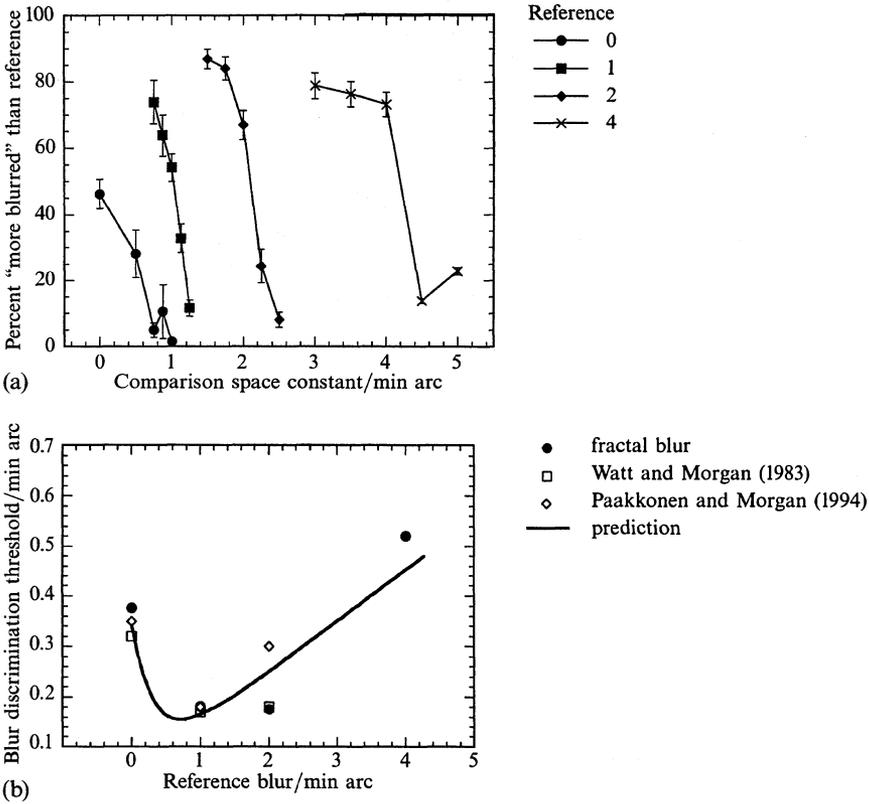


Figure 6. (a) Psychometric functions for blur discrimination in random fractal patterns, for four different reference blurs. Each data point is the mean (± 1 SE) of four naive observers, each making 40 judgments. (b) Blur discrimination thresholds calculated from the data in (a) with the use of logistic regression (circles). For comparison, blur discrimination thresholds for single intensity edges, reported by Watt and Morgan (1983) and Paakkonen and Morgan (1994) are also shown. The continuous line through the data represents predictions from the frequency-based model of blur discrimination, described in detail in the text.

thresholds for Gaussian blur in single edges, replotted from Watt and Morgan (1983, figure 3), and from Paakkonen and Morgan (1994, figure 3). There is good agreement between the studies, despite the use of very different stimuli. Thresholds are in the region of 0.2 to 0.6 min.

All the blur discrimination thresholds plotted in figure 6b show an initial improvement in discrimination at small reference blurs. So-called 'dipper' functions have also been reported for contrast discrimination (eg Nachmias and Sansbury 1974; Legge and Foley 1980), and for spatial discrimination tasks (eg Westheimer and McKee 1977). Can the dipper effect be accommodated by the frequency-based scheme outlined earlier? According to this scheme, blur discrimination should depend on detection of small changes in the highest spatial frequency present in the pattern. Smallman et al (1996) have recently shown that the Weber fraction for spatial-frequency discrimination is fairly constant at about 10% at high frequencies. As figures 4 and 5 show, the relationship between the highest spatial frequency remaining in a blurred pattern and blur space constant is well-described by a simple power function:

$$b = \frac{32}{s}, \quad (3)$$

where s is spatial frequency and b is blur space constant. It is therefore possible to

derive predictions for blur discrimination on the basis of just-discriminable differences in the highest frequency in the blurred patterns. Frequency discrimination thresholds reported in Smallman et al collapsed above approximately 45 cycles deg^{-1} , as stimuli approached the resolution limit of the visual system. So according to the power-law relationship in equation (3), we can assume that even unfiltered patterns (reference blur equal to zero) are subject to an intrinsic Gaussian filter with a space constant of 0.71 min. This value is in good agreement with estimates of the point-spread function of the eye (eg Campbell and Gubisch 1966; Gubisch 1967). When additional extrinsic blur is applied to the image by convolution with a Gaussian filter, then according to the additivity of variances rule for convolution the total blur b_t present in the image is given by

$$b_t = [(b_e^2 + b_i^2)]^{1/2}, \quad (4)$$

where b_i is the intrinsic blur space constant, and b_e is the extrinsic blur space constant. A just-discriminable (ie 10%) change in highest image frequency from 45 to 40.5 cycles deg^{-1} corresponds [equation (3)] to a change in total image blur from 0.71 to 0.79 min. According to the relationship in equation (4), with intrinsic blur taken as 0.71 min, the extrinsic blur space constant required to effect this increase in total blur is 0.34 min. This represents the predicted blur discrimination threshold from the frequency-based model, at a reference blur of zero. It is in good agreement with the data in figure 6b. Predictions were calculated for other reference blurs in the same way, by assuming a constant Weber fraction of 10% for high-frequency discrimination, and are plotted as the continuous curve in figure 6b. The frequency-based scheme captures the dipper function very well, and also predicts the magnitude of blur discrimination thresholds. The dipper effect occurs because, at near-zero reference blurs, relatively large extrinsic blur space constants are required to overcome the contribution of the intrinsic blurring filter. As reference blur increases beyond the size of the intrinsic filter, the latter's contribution diminishes. The reference blur at which the dipper function reaches its minimum can be taken as an estimate of the intrinsic filter space constant (see Watt 1988).

Is it possible to relate the blur discrimination thresholds in figure 6 to Tadmor and Tolhurst's (1994) spectral slope discrimination data? The amplitude spectra of Tadmor and Tolhurst's stimuli were linear, whereas blur introduces an inflexion point in the amplitude spectrum of random fractal patterns. The location of the inflexion point on the frequency axis depends on the blur space constant. Frequency components below it decline in amplitude according to the $1/f$ rule, but amplitudes above it fall much more rapidly. It seems implausible to suggest that the visual system finds the best-fitting straight line through the whole spectrum, and uses its slope to estimate blur extent. It is more likely that the system can discriminate changes in the location of the point of inflexion on the basis of changes in the response of high-frequency-tuned channels, as discussed above. Tolhurst and Tadmor (1997) proposed that discrimination of changes in slope can be modelled as discrimination of changes in the local contrast signalled by spatial-frequency-tuned channels in the visual system. Tadmor and Tolhurst (1994) reported that the subjective appearance of images with high spectral-slope values was akin to blur. Tolhurst and Tadmor (1997) showed that slope discrimination in these stimuli depended on responses in high-spatial-frequency channels. To the extent that slope discrimination relies on the same processes as blur discrimination, this account is consistent with the frequency-based scheme for blur discrimination outlined earlier.

5 General discussion and conclusions

It is interesting to compare the psychophysical data with computed values of blur and depth cues plotted in figure 2. According to the discrimination thresholds in figure 6, blur information is not available over distances relatively close to fixation. For accommodation at 1 m, figure 2a predicts that only objects farther than about 1.3 m (or nearer than 0.7 m) will be sufficiently blurred to provide a reliable depth cue. This prediction agrees well with Campbell's (1957) psychophysical estimate of the eye's depth of field. Stereopsis operates over binocular disparities covering a range of approximately 0.5 to 40 min (Schor and Wood 1983). Applied to the same fixation distance (figure 2a), this would indicate that precise stereo information is available only up to distances of about 1.3 m. The picture is somewhat different for accommodation at 4 m (figure 2b). The stereo cue should be effective up to a range of about 25 m, and the blur cue should be available only over distances beyond 25 m, though the blur cue is already approaching an asymptotic value at this distance. It therefore seems plausible that the two depth cues operate in a complementary fashion, with stereo covering relatively near distances and blur covering relatively far distances. The tendency for the two cues to reach asymptote at large distances has little effect on the effectiveness of the stereo cue but limits the precision of the blur cue, because the latter only operates over larger distances. Blur may therefore provide only qualitative information at such distances.

The argument presented here is obviously rather simplistic, because it neglects the effect of changes in pupil diameter, which re-scale all blur values in the image (figure 3). However, it is worth noting that there is a well-known 'near-reflex' change in pupil diameter. Convergence of the eyes and accommodation on a near object are usually accompanied by pupillary contraction (eg Thompson 1981), so the relationship between stereopsis and pupil diameter is not entirely arbitrary.

As the earlier discussion made clear, estimates of image blur could be provided in a number of ways. Schemes based on multiple images seem unlikely candidates for human vision. Space-domain schemes in the psychophysical literature were developed in the context of 1-D step intensity profiles, and it is not clear how well they generalise to complex 2-D scenes. The frequency-domain scheme seems well-suited to the statistics of natural images, and offers a good account of blur discrimination data. However, to understand more fully the importance of blur as a depth cue it is important to study it specifically in the context of depth perception. Such studies are currently underway in our laboratory.

Acknowledgements. This research was supported by a grant from the Engineering and Physical Sciences Research Council. I thank Linda Murdoch for assistance in data collection.

References

- Berkson J, 1953 "A statistically precise and relatively simple method of estimating the bio-assay with quantal response based on the logistic function" *Journal of the American Statistical Association* **48** 565–599
- Blakemore C, Nachmias J, Sutton P, 1970 "Spatial frequency shift: evidence for frequency-selective neurones in the human brain" *Journal of Physiology (London)* **210** 727–750
- Brown J, Weisstein N, 1988 "A spatial frequency effect on perceived depth" *Perception & Psychophysics* **44** 157–166
- Campbell F, 1957 "Depth of field of the human eye" *Optica Acta* **4** 157–164
- Campbell F, Gubisch R, 1966 "Optical quality of the human eye" *Journal of Physiology (London)* **186** 558–578
- Campbell F, Nachmias J, Jukes J, 1970 "Spatial frequency discrimination in human vision" *Journal of the Optical Society of America* **60** 555–559
- Campbell F, Robson J, Westheimer G, 1959 "Fluctuations of accommodation under steady viewing conditions" *Journal of Physiology (London)* **145** 579–594
- De Valois R, Albrecht D, Thorell L, 1982 "Spatial frequency selectivity of cells in macaque visual cortex" *Vision Research* **22** 545–559

- De Valois R, De Valois K, 1988 *Spatial Vision* (New York: Oxford University Press)
- Ens J, Lawrence P, 1993 "An investigation of methods for determining depth from focus" *IEEE Transactions on Pattern Analysis and Machine Intelligence* **15** 97–107
- Field D, 1987 "Relations between the statistics of natural images and the response properties of cortical cells" *Journal of the Optical Society of America A* **4** 2379–2394
- Georgeson M, 1994 "From filters to features: location, orientation, contrast and blur", in *Higher Order Processing in the Visual System* Ciba Foundation Symposium 184 (Chichester: John Wiley) pp 147–165
- Georgeson M, Harris M, 1984 "Spatial selectivity of contrast adaptation: models and data" *Vision Research* **24** 729–741
- Gubisch R, 1967 "Optical performance of the human eye" *Journal of the Optical Society of America* **57** 407–415
- Howard I P, Rogers B J, 1995 *Binocular Vision and Stereopsis* (Oxford: Oxford University Press)
- Klein S, Stromeyer C, Ganz L, 1974 "The simultaneous spatial frequency shift: a dissociation between the detection and perception of gratings" *Vision Research* **14** 1421–1432
- Knill D, Field D, Kersten D, 1990 "Human discrimination of fractal images" *Journal of the Optical Society of America A* **7** 1113–1123
- Kotulak J, Schor C, 1986 "Temporal variations in accommodation during steady-state conditions" *Journal of the Optical Society of America A* **3** 223–227
- Legge G, Foley J, 1980 "Contrast masking in human vision" *Journal of the Optical Society of America* **70** 1458–1471
- Marc R, Sperling H, 1977 "Chromatic organisation of primate cones" *Science* **196** 454–456
- 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
- Nachmias J, Sansbury R, 1974 "Grating contrast: discrimination may be better than detection" *Vision Research* **14** 1039–1042
- Paakkonen A, Morgan M, 1994 "Effects of motion on blur discrimination" *Journal of the Optical Society of America A* **11** 993–1002
- Peitgen H, Saupe D (Eds), 1988 *The Science of Fractal Images* (Berlin: Springer)
- Pentland A, 1987 "A new sense for depth of field" *IEEE Transactions on Pattern Analysis and Machine Intelligence* **9** 523–531
- Schaaf A van der, Hateren J van, 1996 "Modelling the power spectra of natural images: statistics and information" *Vision Research* **36** 2759–2770
- Schor C, Wood I, 1983 "Disparity range for local stereopsis as a function of luminance spatial frequency" *Vision Research* **23** 1649–1654
- Smallman H S, MacLeod D I A, He S, Kentridge R W, 1996 "Fine grain of the neural representation of human spatial vision" *Journal of Neuroscience* **16** 1852–1859
- Subbarao M, Gurumoorthy N, 1988 "Depth recovery from blurred edges", in *Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition* (Washington, DC: IEEE Computer Society) pp 498–503
- Tadmor Y, Tolhurst D, 1994 "Discrimination of changes in the second-order statistics of natural and synthetic images" *Vision Research* **34** 541–554
- Thompson H, 1981 "The pupil", in *Adler's Physiology of the Eye* Ed. R Moses (St Louis, MN: Mosby) pp 326–356
- Tolhurst D, Tadmor Y, 1997 "Discrimination of changes in the slopes of the amplitude spectra of natural images: band-limited contrast and psychometric functions" *Perception* **26** 1011–1025
- Watt R J, 1988 *Visual Processing* (Hove, Sussex: Lawrence Erlbaum Associates)
- Watt R, Morgan M, 1983 "The recognition and representation of edge blur: evidence for spatial primitives in human vision" *Vision Research* **23** 1465–1477
- Westheimer G, McKee S, 1977 "Spatial configurations for hyperacuity" *Vision Research* **17** 941–947
- Winn B, Charman W, Pugh J, Heron G, Eadie A, 1989 "Perceptual detectability of ocular accommodation microfluctuations" *Journal of the Optical Society of America A* **6** 459–462