Spatial Vision, Vol. 9, No. 1, pp. 151–161 (1995)

Human discrimination of surface slant in fractal and related textured images

P. J. PASSMORE* and A. JOHNSTON

Department of Psychology, University College London, Gower St., London, UK

Received 2 May 1994; revised 5 December 1994; accepted 6 December 1994

Abstract—Slant-discrimination thresholds were measured for textures with the property that their power spectra, when log transformed, are inversely proportional to log spatial frequency: $P(f) \propto f^{-\beta}$. As the exponent β changes from high values to low values, the slope of the power spectrum of the image decreases. As the parameter passes through values in the fractal range the resulting texture changes from having the appearance of a cloud-like surface through to a granite-like surface. Exponents below the fractal range produce textures that converge towards the appearance of a random grey-level noise pattern. Since fractal patterns are self-similar at a range of scales, one might think it would be difficult to recover changes in depth in fractal images; however, slant-discrimination thresholds did not differ substantially as a function of the slope of the power spectrum. Reducing the size of the viewing aperture increased thresholds significantly, suggesting that slant discrimination benefits from a global analysis. The effect of texture regularity on perceived slant was investigated using bandpassed fractal textures. As the bandwidth of a bandpass filter is reduced, the bandpassed texture was perceived to be increasingly more slanted than its fractal counterpart.

1. INTRODUCTION

Symmetries involve transformations of an object that leave the appearance of the object unchanged. The mathematical ideal can be relaxed to provide the concept of a perceptual symmetry, which we can define as a transformation of a pattern that, to the human observer, leaves the pattern indistinguishable from the original, in the sense that they cannot be distinguished without detailed scrutiny. For example, two patches of the same texture may appear interchangeable or two images of the same fractal surface, produced at different scales, may appear to be self-similar.

Surface texture has, since Gibson (1950), been seen as an important cue to the orientation of surfaces. The two assumptions mostly commonly made about textures in shape-from-texture algorithms are that the surface texture is either isotropic (Witkin, 1981; Blake and Marinos, 1990; Brown and Shvaytser, 1990), or homogeneous (Aliomonos, 1988; Kanatani and Chou, 1989). Isotropy implies that there is

*To whom correspondence should be addressed. Present address: School of Computing, University of North London, 166-220 Holloway Rd., London N7, UK.

no orientational bias in the texture while homogeneity embodies the constraint proposed by Gibson (1979) that there are equal amounts of texture per unit of surface area. Shape-from-texture algorithms typically estimate slant by attempting to apply an inverse perspective or orthographic transform to the image texture to maximize isotropy or homogeneity in the resulting pattern.

Psychophysical studies of shape-from-texture have demonstrated the importance of a number of cues affecting perceived slant. Gibson (1950) found, using wallpaper textures and a cross-modal matching task, that perception of slant from texture is improved by increasing the regularity of textures. More recently, investigation has focused on the relative importance of three texture gradients: perspective, compression, and density. If we consider a uniform texture in the ground plane, the perspective gradient is the x-axis width of texture elements (texels) projected onto the image plane orthogonal to the line of sight, the compression gradient is the ratio of y/x axis measures on the projected plane, and the density gradient corresponds to the number of texels per unit visual angle.

A number of studies have suggested that the compression gradient is the most important for shape-from-texture (Cutting and Millard, 1984; Todd and Akerstrom, 1987; Cumming *et al.*, 1993). Cumming *et al.* proposed that shape-from-texture involved assumptions of isotropy rather than homogeneity as they found that disrupting isotropy reduced the relative influence of texture in judgments of surface shape in stereoscopic displays. Stone (1993) suggests that the critical factor is that the distribution of the orientation of the local tangent vectors to the brightness contours in the image is invariant with position — a more constrained form of anisotropy. There has also been some debate whether shape-from-texture is a local or global process, with Stevens (1981) suggesting that the process may be considered local under the assumption of approximately circularly symmetric texture elements, whilst Cumming *et al.* (1993) find evidence for global processes. Elongated texture elements, which formed elliptic contours on the test surfaces, were not interpreted as regular circular contours in an alternative depth plane.

Psychophysical investigations of shape-from-texture often use idealized textures which are typically produced by tiling a planar or curved surface with standardized elements such as lines or circles (e.g. Cutting and Millard, 1984) or carving from volumetric textures containing ideal shapes such as spheres and ellipsoids (Cumming *et al.*, 1993). Fractal images, on the other hand, have the appearance of naturally occurring textures (such as those produced by clouds or granite) and many images of different natural scenes have been shown to have amplitude spectra in the range of the fractal model (Field, 1987). A fractal surface is characterized by its fractal dimension, which corresponds to the roughness of the surface, a property that is invariant over scale.

When a textured surface is slanted, there is a change in scale in the image with distance as the more distant surface patches are projected smaller in the image. As a fractal texture has spatial structure which is invariant with scale, it has been assumed that surface orientation would be difficult to recover from fractal textures. However, as Stevens (1981) points out, although in the direction orthogonal to the tilt of the surface perspective projection simply results in a spatial scaling of the image features, in the direction in which the texture is slanted there is an anisotropic transformation (foreshortening) which varies as a function of spatial position. If subjects are sensitive to this property of perspective projection then it should be possible to see depth in slanted fractal textures. We decided to test this idea directly by measuring slant-discrimination thresholds for fractal textures.

2. GENERAL METHODS

2.1. Stimulus generation

Many natural images have a $1/f^2$ power spectrum (Field, 1987). Note that a $1/f^2$ power spectrum corresponds to a $1/f$ amplitude spectrum. We generated synthetic textures with the property that their power spectra, when log transformed, are inversely proportional to log spatial frequency:

$$P(f) \propto f^{-\beta}.$$

As the exponent β changes from high values to low values, the slope of the power spectrum of the image decreases. The resulting texture changes from having the appearance of a cloud-like surface through a granite-like texture to a random grey-level noise pattern. Knill *et al.* (1990) explain that, for a fractal model based on fractional Brownian motion, there is a relationship between the fractal dimension and the slope of the power spectrum. For brightness surfaces the exponent β can be shown to be related to the fractal dimension D by:

$$\beta = 8 - 2D,$$

where D falls in the range 2 to 3 (Knill *et al.*, 1990). Thus a fractal image would give rise to values of the exponent β between 2 and 4. We can of course design images with power spectra which have values of β outside this range. These patterns would not be representative of the class of images that can be generated by a process based on fractional Brownian motion.

The spectral synthesis or Fourier filtering method was used to generate stimuli (Saupe, 1988). The stimuli were constructed in frequency space with a given slope specified by the square root of the β exponent and a random phase spectrum. An inverse Fourier transform was then applied to produce 2048 by 2048 pixel images. The images were rotated in 3D space to produce 512 by 512 pixel texture gradients. Aliasing was reduced by supersampling and perspective averaging. The intensity of a pixel in the frontal plane was determined by averaging over an elliptic area of the rotated surface that projects onto the neighbourhood of that pixel. The average value was computed by projecting an elliptic grid of points onto the rotated surface and summing over the intensity values of those points, which are calculated by nearest neighbour grey-level bilinear interpolation. The resulting stimuli were scaled to have the same Michelson contrast. The images were displayed with eight-bit precision on

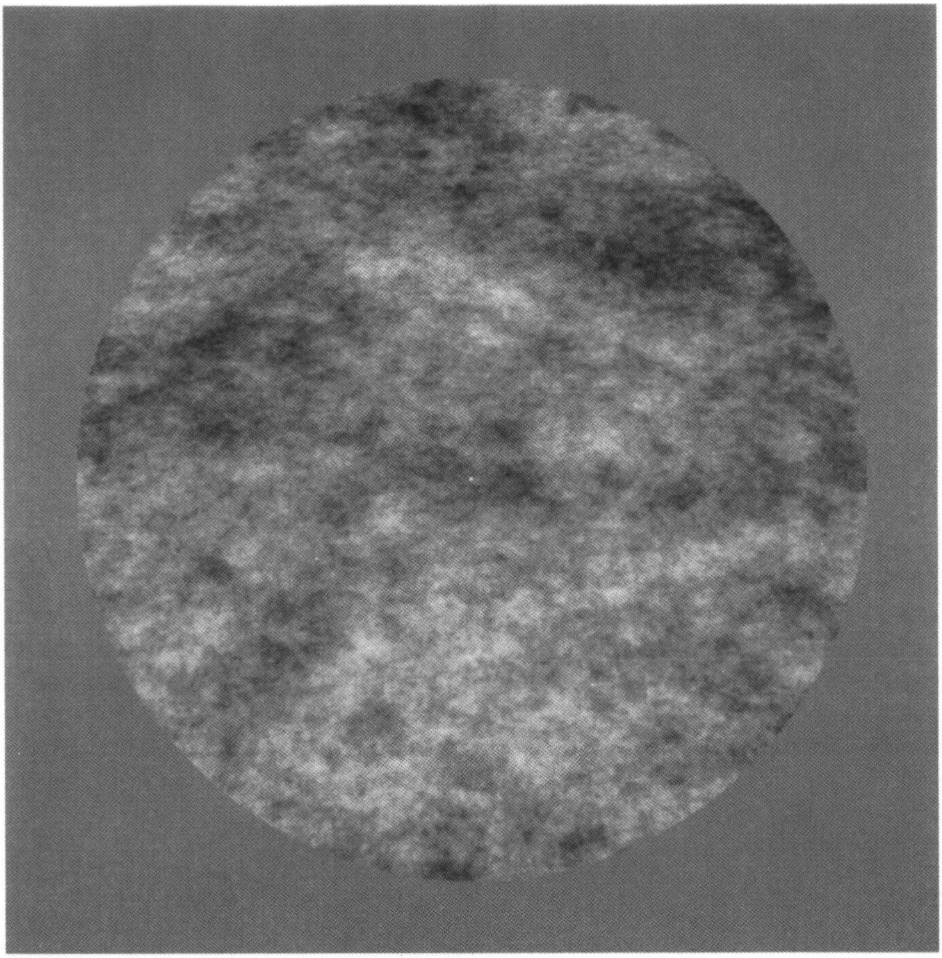

Figure 1. This is an example fractal stimulus ($\beta = 2.2$) slanted by 50 deg towards the ground plane.

a nineteen-inch Sony Trinitron monitor under the control of a Sun SPARCstation 330. In order to linearize the display, a lookup table of luminance values was determined with a micro-photometer and used to control stimulus brightness. Figure 1 shows an example fractal stimulus.

2.2. *General procedure*

Slant-discrimination thresholds were measured by a method of constants for stimuli that depicted textures that had been rotated around a central horizontal axis; subjects were required to decide whether the stimuli were slanted towards the ground plane or the sky plane. No feedback was given during the experiments but subjects were required to conduct a few practice runs. The textures were occluded by a circular window, which subtended 27 deg of visual angle. The size of the aperture was randomly jittered by up to ±12% between presentations, and the position of the aperture in the window was randomly jittered by up to ±1.7 deg in a random direction between trials. These randomizations were employed to prohibit subjects using simple comparisons of texture size over trials as a cue, thus forcing them to rely on estimates of the slant of the surface. Stimuli were viewed monocularly from 25 cm. The head was kept still using a headrest, and subjects were required to fixate a central point embedded in the image. Due to the technical difficulty of generating completely novel stimuli from trial to trial, stimuli were chosen from a pool of 36 images grouped into 9 deg of slant. The stimuli in each group had identical amplitude spectra but different random phase spectra. Thresholds were calculated by probit analysis (Finney, 1971) and each threshold was the result of at least 280 judgments. Discrimination threshold was defined as the standard deviation of the error distribution and corresponds to the 84% point on the psychometric function.

3. EXPERIMENT 1. SLANT DISCRIMINATION THRESHOLDS AS A FUNCTION OF THE SLOPE OF THE POWER SPECTRUM

Slant discrimination was measured as a function of the slope of the power spectrum with the exponent β ranging from 0.2 to 3.8. This includes stimuli whose β exponent fell within the fractal range and extends outside it to include patterns whose power spectrums have shallower slopes. The slants used were in the range of 0 to 50 deg in either direction with respect to the image plane. The two authors and two naive subjects, one of whom (WC) is stereo blind due to a late-corrected strabismus, participated in the experiment. It should be noted that WC's results did not differ substantially from those of other subjects in these experiments.

Figure 2 shows the slant discrimination thresholds averaged over the four subjects as a function of the exponent β. We found that subjects could see surface slant in these displays. Thresholds were slightly higher for the fractal range than for the textures with flatter power spectra but there were no substantial differences. Thresholds ranged from around 20 to 35 deg increasing slightly as the slope of the amplitude spectrum

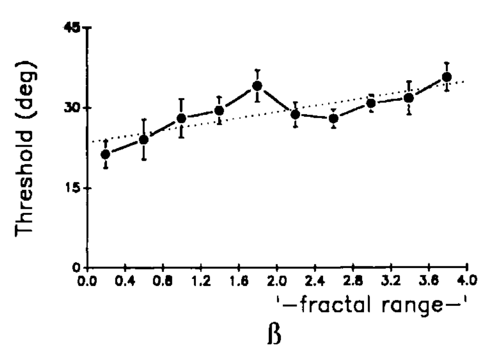

Figure 2. Averaged slant discrimination thresholds are plotted as a function of the slope of the power spectrum (β) for four subjects. Although there is a slight decrease in threshold with decrease in the slope the effect is small and demonstrates a substantial degree of scale invariance for slant discrimination with these textures.

increased. There was some variation between individuals but no radical effect of the slope of the amplitude spectrum on discrimination thresholds was found.

4. EXPERIMENT 2. SLANT DISCRIMINATION THRESHOLDS: EFFECTS OF FILTERING GAUSSIAN NOISE AND FLAT SPECTRA WITH A $f^{-\beta}$ FILTER

In Experiment 1 we used the spectral synthesis method to generate fractal textures. In essence this technique involves filtering a Gaussian random variable by a $f^{-\beta}$ spectrum. The amplitude of each spectral component is drawn from a Gaussian distribution and then multiplied by $f^{-\beta}$. However, the expected power spectrum for an ensemble of fractal images with a particular fractal dimension would be a smooth sloping line, the slope of which is related to the fractal dimension of the ensemble. In order to quantify the effects of the random variable employed in Experiment 1 we compared slant discrimination thresholds for Gaussian noise and flat spectra both of which were modified by a $f^{-\beta}$ filter. The experimental procedure was the same as in Experiment 1. Two subjects, the authors, participated in this experiment.

The results are shown in Figs 3a and b. While thresholds for subject PP are generally lower than for subject AJ, both subjects have similar results for both the Gaussian and flat spectra conditions, demonstrating that the random variation in the amplitude of the power spectra in Experiment 1 had no radical effects on slant perception.

5. EXPERIMENT 3. EFFECT OF APERTURE SIZE ON SLANT DISCRIMINATION THRESHOLDS

It is widely accepted that there exists a range of spatial channels at each retinal position in the human visual system, and that the range of spatial analysis shifts progressively to a coarser spatial scale moving towards the periphery of the visual field. Increasing the retinal image size of a slanted stimulus will increase the number and scale of spatial filters that the visual system may use to process a stimulus. This may improve performance. If the visual system can determine slant on the basis of

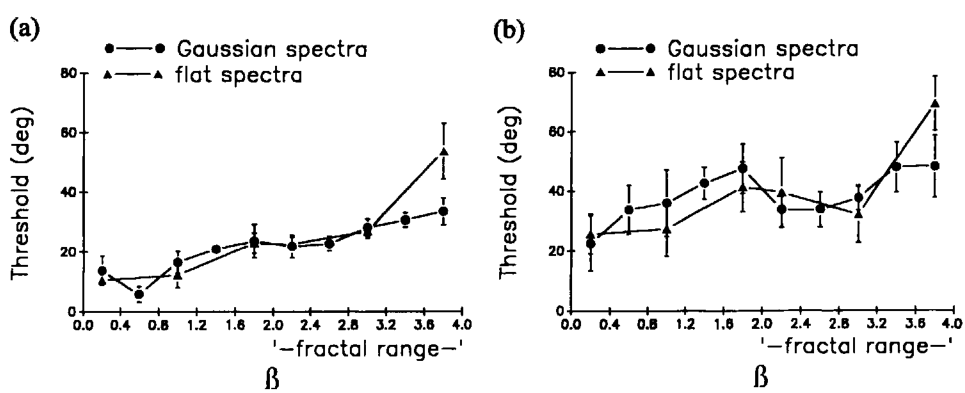

Figure 3. Slant discrimination thresholds plotted as a function of the slope of the power spectrum for Gaussian noise and flat power spectra. (a) subject PP, (b) subject AJ.

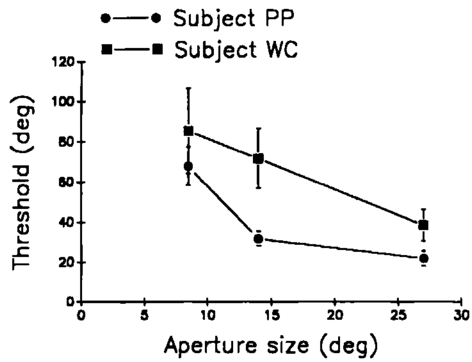

Figure 4. Slant discrimination thresholds are plotted as a function of aperture size. Thresholds increase with a decrease in aperture size suggesting that slant discrimination is made on the basis of global information.

local information then we would expect that the reduction of stimulus aperture size would not greatly affect slant discrimination thresholds. To investigate whether slant discrimination from texture is achieved by local or global processes, we used the same experimental procedures as in Experiment 1 but varied the aperture size. The power spectra of the stimuli were fixed ($\beta = 2.2$). The apertures subtended 27, 14, and 8.5 deg of visual angle. One author (PP) and a naive subject (WC) participated in the experiment.

Figure 4 shows slant discrimination thresholds as a function of aperture size. For both subjects thresholds increased significantly with a decrease in aperture size. Thus, for these textures, slant discrimination benefits from access to global information.

6. EXPERIMENT 4. SLANT DISCRIMINATION OF BANDPASSED FRACTAL TEXTURES

For fractal textures, bandpass filtering tends to increase regularity. The regularity increases systematically as bandwidth is reduced. For example, when the exponent β has the value 2.2, the texture changes from being granite-like to being more mottled in appearance with local blob-like texels becoming discriminable in the bandpassed texture. Increasing regularity can also be considered to introduce a greater degree of perceptual symmetry in a pattern. In the one-dimensional case, filtering a noise function with a sufficiently narrow bandpass filter will produce a single sine wave which has reflective, translational, and rotational symmetries.

In this experiment, the effect of texture regularity on perceived slant was investigated by bandpass filtering fractal textures prior to 3D rotation. We used a range of bandwidths centred on a frequency of 512 cycles per image. The image size was 2048 by 2048 pixels. Perceived difference in slant between the original fractal stimuli and bandpassed fractal stimuli was measured using a two-alternative forced-choice task. The fractal stimuli were randomly chosen from a pool of four images that had the same exponent ($\beta = 2.2$) but different random phase spectra, and were slanted at 50 deg towards ground plane. The exponent value of $\beta = 2.2$ was arbitrarily chosen, being near the middle of the range of values used but in the fractal range. Bandpassed

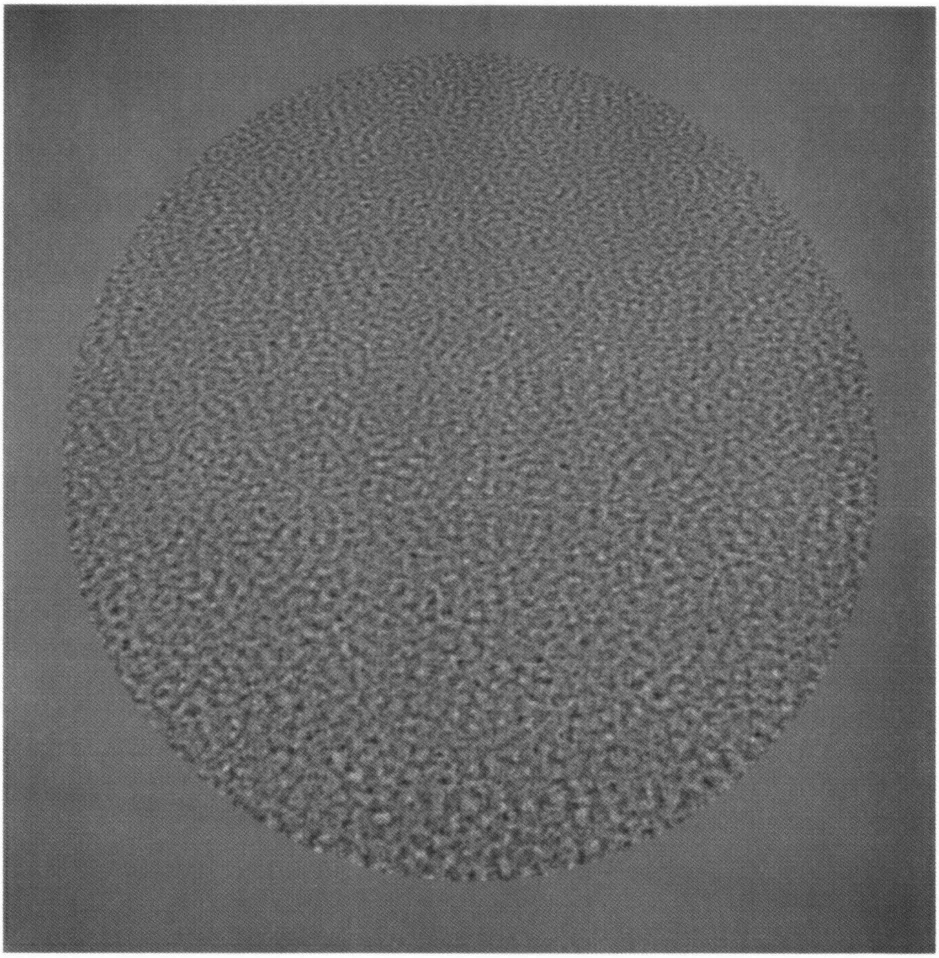

Figure 5. This is an example stimulus used in Experiment 4. It was generated by bandpass filtering a fractal image prior to rotation using a filter with a bandwidth of 796 cycles per image. Note that this texture, like that of Fig. 1, is also slanted by 50 deg towards the ground plane but appears to most observers to be more slanted owing to the effects of texture regularity. We may consider the bandpass process as increasing the symmetry of the texture.

stimuli were chosen from a pool of images in which nine levels of slant (from 25 to 50 deg) and four random phase spectra were combined. The experimental procedure followed that used in Experiment 1 except that the fractal and bandpassed texture were presented consecutively. For each trial the order of presentation was randomized. The subjects' task was to indicate which texture appeared more slanted. Stimuli were displayed for approximately 2 sec and an inter-stimulus grey-level image was displayed for approximately 3 sec between the two textures. An example bandpassed stimulus is shown in Fig. 5. Two subjects, the authors, participated in the experiment.

Figure 6 shows the perceived difference in slant as a function of bandwidth. For both subjects, perceived difference in slant increases with the decrease of bandwidth.

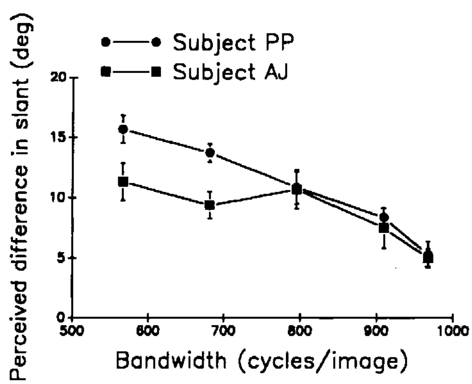

Figure 6. Perceived difference in slant is plotted as a function of the bandwidth of the bandpass filter. The increase in perceived difference in slant for both subjects with decrease in bandwidth size is due to the effects of texture regularity.

This replicates, for fractal textures, the finding that as texture regularity increases slanted surfaces appear more slanted. This result agrees with the findings of Gibson (1950), Flock and Moscatelli (1964) and Phillips (1970) that increases in the regularity of texture increases perceived slant. Applying a bandpass filter to the texture makes it visually segmentable. Although self-similarity over changes in scale is reduced, there is an increase in self-similarity with respect to spatial position, since translation involving multiples of the spatial period of the pattern would not be easily discriminable.

7. DISCUSSION

Because fractal surfaces display self-similarity under changes in scale, it is often thought that it would be difficult to see a change of distance in a fractal surface since changing distance simply results in a change of scale in the retinal image. However, we have shown that subjects can make judgments of surface slant for surfaces defined by fractal textures and the slant discrimination thresholds are similar to those for textures with flatter power spectra, which appear rather like random grey-level textures. The results of Experiment 1 show that there is little effect of the slope of the amplitude spectrum (and consequently fractal dimension) on slant discrimination thresholds. Although we might expect difficulties in recovering changes in local scale brought about by changes in slant, there are a number of possible sources of information available in slanted fractal textures that might be used to recover the slant. Any effects must be related to perspective projection since our subjects made decisions about slant for a range of stimuli that included surface rotations in both directions with respect to the vertical. The scaling effects resulting from orthographic projection are identical for surfaces slanted by the same amount in either direction. Subjects therefore must be using perspective cues. One possible source of information is in the distribution of local orientations along isophote contours, or lines of equal reflectance (Witkin, 1981; Stone, 1993). When an isotropic textured surface is slanted, there are local changes in the distribution of orientations that can be used to recover slant. This effect results from the same geometry that gives rise to linear perspective cues, which also signal depth through the orientation of image contours.

For an ideal texture produced by tiling a plane with circles, it may be possible to base slant discrimination on local estimates as proposed by Stevens (1981). However, a characteristic of the textures used in the first three experiments is that they are not visually segmentable into local texture elements. Thus the experiments provide further evidence for the role of global processes as against local processes in slant perception, since slant discrimination thresholds are markedly increased for small aperture sizes. This effect may also provide an explanation for the increases in thresholds for the reduced aperture conditions if we accept that reducing aperture size reduced the ability of subjects to make use of foreshortening and convergence cues.

The techniques used in Experiment 4 allowed us to systematically manipulate texture regularity by varying the pass-band of the fractal grey-level texture and we have replicated earlier findings that regular textures appear to be more slanted. We can characterize this manipulation in terms of increasing the symmetry in the image; as an image is bandpass filtered with increasing narrow spatial filters it becomes increasingly more cyclical, visually segmentable, and regular. Thus, for a fractal image, bandpass filtering decreases self-similarity over scale and increases self-similarity over spatial position. Both of these manipulations can be expected to act together to improve slant discrimination.

Acknowledgement

This work was supported by a grant from the SERC Image Interpretation Initiative.

REFERENCES

Aloimonos, J. (1988). Shape from texture. *Biol. Cybernet.* **58**, 345–360.

Blake, A. and Marinos, C. (1990). Shape from texture: estimation, isotropy and moments. *Artif. Intel.* **45**, 323–380.

Brown, L. G. and Shvaytser, H. (1990). Surface orientation from projective foreshortening of isotropic texture autocorrelation. *IEEE Trans. PAMI* **12** (6), 584–588.

Cumming, B. G., Johnston, E. B. and Parker, A. J. (1993). Effects of different texture cues on curved surfaces viewed stereoscopically. *Vision Res.* **33**, 827–838.

Cutting, J. E. and Millard, R. T. (1984). Three gradients and the perception of flat and curved surfaces. *J. Exp. Psychol. Gen.* **113**, 198–216.

Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A* **4**, 2379–2394.

Finney, D. J. (1971). *Probit Analysis*, 3rd edn. Cambridge University Press, Cambridge, UK.

Flock, H. R. and Moscatelli, A. (1964). Variables of surface texture and accuracy of space perceptions. *Percept. Motor Skills* **19**, 327–334.

Gibson, J. J. (1950). *The Perception of the Visual World.* Houghton Mifflin, Boston.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception.* Houghton Mifflin, Boston.

Kanatani, K. and Chou, T. (1989). Shape from texture: general principle. *Artif. Intel.* **38**, 1–48.

Knill, D. C., Field, D. and Kersten, D. (1990). Human discrimination of fractal images. *J. Opt. Soc. Am. A* **7** (6), 1113–1123.

Phillips, R. J. (1970). Stationary visual texture and the estimation of slant angle. *Q. J. Exp. Psychol.* **22**, 389–397.

Saupe, D. (1988). Algorithms for random fractals. In: *The Science of Fractal Images*. H. Peitgen and D. Saupe (Eds). Springer-Verlag, New York, pp. 71–136.

Stevens, K. (1981). The information content of texture gradients. *Biol. Cybernet.* **42**, 95–105.

Stone, J. V. (1993). Shape from local and global analysis of texture. *Phil. Trans. Roy. Soc. Lond. B* **339**, 53–65.

Todd, J. T. and Akerstrom, R. A. (1987). Perception of three-dimensional form from patterns of optical texture. *J. Exp. Psychol: Human Percept. Perform.* **13**, 242–255.

Witkin, A. P. (1981). Recovering surface shape and orientation from texture. *Artif. Intel.* **17**, 17–47.