USING FREQUENCY- AND ORIENTATION- TUNED CHANNELS TO DETERMINE SURFACE SLANT

PAUL KUBE

The Institute of Cognitive Studies
and
The Department of Electrical Engineering and Computer Sciences
University of California, Berkeley

INTRODUCTION

The thesis that spatial frequency analysis of the retinal image plays a fundamental role in the early stages of visual processing has received considerable evidential support recently, both from psychophysics (e.g., (Campbell & Robson 1968), (Sachs, Nachmias & Robson 1971), (DeValois & DeValois 1980), (Daugman 1984a)) and neurophysiology (e.g., (Campbell et al. 1969), (Maffei & Fiorentini 1973), (Schiller, Finlay & Volman 1976), (DeValois, Yund & Hepler 1982)). What role the spatial frequency information so extracted plays in later stages of processing remains an open question, though one that has increasingly received the attention of theoreticians and modellers (e.g., (Ginsburg 1978), (Watson 1983), (Janez 1983), (Daugman 1984b)). Here we present results which show how this information could be used in the determination of the orientation of environmental surfaces, and discuss the limits of such a model as an account of human perceptual capacities.

ANISOTROPY, SLANT, AND SPECTRUM

Suppose a planar environmental surface is oriented perpendicular to the line of sight and produces an irradiance pattern $I(x,y)$ on the retina, relative to some suitably chosen retinal coordinate frame. Rotating the environmental surface about a line parallel to the retinal y axis by an angle σ now foreshortens the retinal image: if the distance to the surface is large enough with respect to the size of the surface for orthographic projection to be a good approximation, the retinal irradiance pattern becomes $I'(x,y)=I(x,y/\cos\sigma)$.[1]

[1] Ignoring effects due to any change in angle between the surface normal and the illuminant vector.

If appropriate assumptions can be made about what the unforeshortened retinal image is like, the slant angle σ can be recovered from the foreshortened image I'. It has been suggested (e.g., by Witkin (Witkin 1981)) that if the retinal image is a foreshortened isotropic image, it should be so interpreted. That is, if $I'(x,y)$ is such that, for some σ, $I'(x,y\cos\sigma)$ is isotropic, then the imaged surface should be supposed to be slanted at angle σ.

Witkin counts an image isotropic if the distribution of tangent directions along contours in the image (or an edge-enhanced version of the image) are approximately uniform on $[0,\pi]$. Here, we suggest that an image be considered isotropic if its Fourier power spectrum is approximately circularly symmetric. This is an acceptable condition for isotropy for the following reason. Consider the definition of the two-dimensional Fourier transform of the retinal image $I(x,y)$:

$$f(\omega,\theta) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-i2\pi(x\omega\cos\theta + y\omega\sin\theta)} I(x,y)\, dy\, dx$$

This expresses the spectrum f in polar coordinates of frequency ω and orientation θ. Fix $\theta=0$. Now

$$\begin{aligned} f(\omega,0) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-i2\pi x\omega}\, I(x,y)\, dy\, dx \\ &= \int_{-\infty}^{\infty} e^{-i2\pi x\omega} \int_{-\infty}^{\infty} I(x,y)\, dy \quad dx \end{aligned}$$

But this is just a *one*-dimensional Fourier transform of the function $I(x)$ such that

$$I(x) = \int_{-\infty}^{\infty} I(x,y)\, dy$$

$I(x)$ represents what $I(x,y)$ is like, 'on average', in the direction parallel to the x axis, and $f(\omega,0)$ just encodes this information as a function of frequency. Now since the orientation of the retinal coordinate frame was arbitrary, it's clear that a 'slice' through the two-dimensional spectrum $f(\omega,\theta)$ for any fixed θ represents what the image is like 'on average' in the direction making angle θ with the x axis; and thus if the spectrum is circularly symmetric, so will the image be. And a circularly symmetric image is intuitively isotropic.

But circular symmetry of images is too strong a necessary condition of isotropy. Requiring only that the Fourier *power* spectrum $F(\omega,\theta) = |f(\omega,\theta)|^2$ be circularly symmetric relaxes the condition in a satisfying way. Now isotropy is insensitive to translations of the retinal coordinate frame (since translation only affects the phase of the spectrum, not its magnitude) and to other perturbations of phase of components of the spectrum. Moreover, the power spectrum represents the foreshortening effects of slant in a systematic way. By the similarity theorem of Fourier analysis (cf. (Bracewell 1978) , p. 244), if $F(\omega,\theta)$ is the power spectrum of the image $I(x,y)$, then

$$F'(\omega,\theta) = \cos^2\sigma\, F\left((\omega^2\cos^2\theta+\omega^2\sin^2\theta\cos^2\sigma)^{1/2},\ \arctan\frac{\sin\theta\cos\sigma}{\cos\theta}\right) \qquad (1)$$

is the power spectrum of $I(x,y/\cos\sigma)$.

The power spectrum characterization of isotropy and Witkin's tangent distribution characterization agree for many images. We have adopted the spectrum characterization because it establishes a relationship between the presence or absence of isotropy in the retinal image and the activation of biologically plausible tuned channels in a way that permits the extraction of slant information from some kinds of images. This relationship will be developed in the remainder of the paper.

THE TUNED CHANNEL MODEL

To say that there are spatial frequency- and oriention-tuned channels in the visual system is just to say that there are neurons in visual cortex whose firing rates are elevated only when the power spectrum of the retinal image shows energy in certain regions of the Fourier plane. The power spectrum of the retinal image can be well represented by the activation level in many such channels.

The point spread functions ('receptive fields') and spectral response of two channels are shown in Figure 1. Channels vary in their nominal orientation and frequency but are assumed to be identical in orientation and relative frequency bandwidth at 30 degrees and one octave respectively. These bandwidth parameters are those used by Jañez (Janez 1983) and are close to the average observed in cortex (DeValois, Yund & Hepler 1982); Watson (Watson 1983) and Daugman (Daugman 1984a) propose slightly different parameters. In the present application, what's important is that the channels be fairly narrowly tuned in frequency and orientation, and that there be enough of them at enough different nominal frequencies and orientations to cover the Fourier plane over the frequency range of interest.

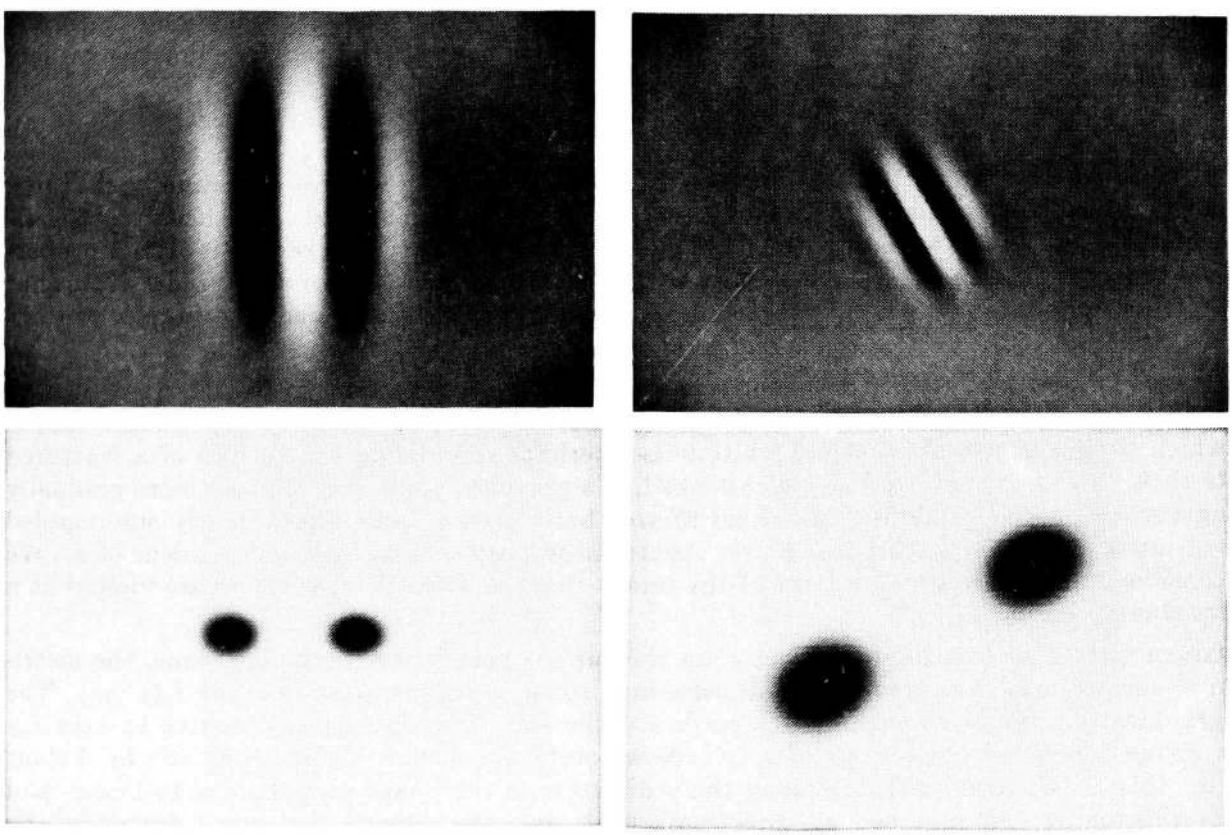

(a) (b)

Figure 1

Point spread (above) and frequency response (below) functions for two tuned channels: (a) frequency ω_0, orientation 0°; (b) frequency $2\omega_0$, orientation 30°. Each channel has half-amplitude bandwidth of one octave in frequency and 30 degrees in orientation.

The activation in a single tuned channel responding to a retinal image is calculated as follows. Let the power spectrum of the image I_i be described by the function of frequency and orientation $F_i(\omega,\theta)$, and that of the response of channel centered at frequency ω_0 and orientation θ_0 be $F_{\omega_0,\theta_0}(\omega,\theta)$; then we say that the activation produced by the image in the channel at f_0,θ_0 is

$$P_{\omega_0,\theta_0}(I_i) = \frac{\int_{\infty 0}^{\infty 2\pi} \int F_i(\omega,\theta)\, F_{\omega_0,\theta_0}(\omega,\theta)\, d\theta\, d\omega}{\int_{\infty 0}^{\infty 2\pi} \int F_{\omega_0,\theta_0}(\omega,\theta)\, d\theta\, d\omega} \qquad (2)$$

That is, channel power is the sum of image frequency power weighted by the response of the channel, and normalized by the total frequency-plane area of the channel.

DETERMINING THE SLANT OF SCALING NOISE SURFACES

The distribution of markings on surfaces in the natural environment is typically random in some sense; highly regular, structured patterns are rare. Here we consider the class of surfaces marked with *isotropic scaling noise*, a class of markings that seem to model some naturally occuring textures.[2] We define such surfaces as those which have spatial radiance functions whose power spectra are of the form

$$F_\beta(\omega,\theta) = k\, \omega^{-\beta}$$

with $\beta \geq 0$, and with random phase. Here β is a parameter that determines how rapidly changing the radiance across the surface is; as β increases, the power at high spatial frequencies on the surface decreases, and so the surface has a more slowly varying pattern marked on it (k is a proportionality constant that can be taken as unity if the markings are suitably normalized). Note that the spectrum is independent of orientation θ, so we should expect such a surface to look statistically the same in every direction (thus *isotropic* scaling noise).

Images of isotropic scaling noise surfaces with varying β are shown in Figure 2(a-d) viewed with zero slant, i.e., 'straight-on', with line of sight normal to the surface. An extreme case is $\beta=0$ which produces two-dimensional white noise, perhaps resembling the surface of a fractured granite rock. $\beta=1.0$ gives a surface which looks like a gravelled walk; $\beta=2.0$ has a more gradually varying texture, a fair rendering of a kind of tree bark; $\beta=3.0$ looks like a gently sun-dappled patch of lawn. As noted, each has a power spectrum proportional to $1/\omega^\beta$, independent of θ. We next consider how the power spectrum of the image changes when these surfaces are viewed at a *nonzero* slant.

Given Cartesian coordinate axes x',y' in the surface parallel to the image plane, the distribution of surface markings is some two-dimensional isotropic scaling noise function $I_\beta(x',y')$. The orthographically projected image of this surface is the same function $I_\beta(x,y)$ relative to axes x,y in the image which are projections of x',y'. Now rotate the surface about its x' axis by a slant angle σ: this causes foreshortening along the y direction in the image proportional to $1/\cos\sigma$, and no foreshortening along the x direction; that is, the image is now described by $I_{\beta,\sigma}(x,y) = I_\beta(x,y/\cos\sigma)$, which is no longer an *isotropic* scaling noise. (Examples of images of

[2] This class is a superset of the surfaces marked with isotropic fractional Brownian plane-to-intensity functions ((Mandelbrot 1983) , ch. 25-27), which restrict β to the range 2 to 4. The imaging of unmarked nonplanar surfaces whose shape is described by an isotropic fractional Brownian plane-to-*elevation* function is another matter, and it doesn't seem to have been yet adequately treated (Pentland's fractal Brownian surfaces (Pentland 1984) are not fractal in elevation).

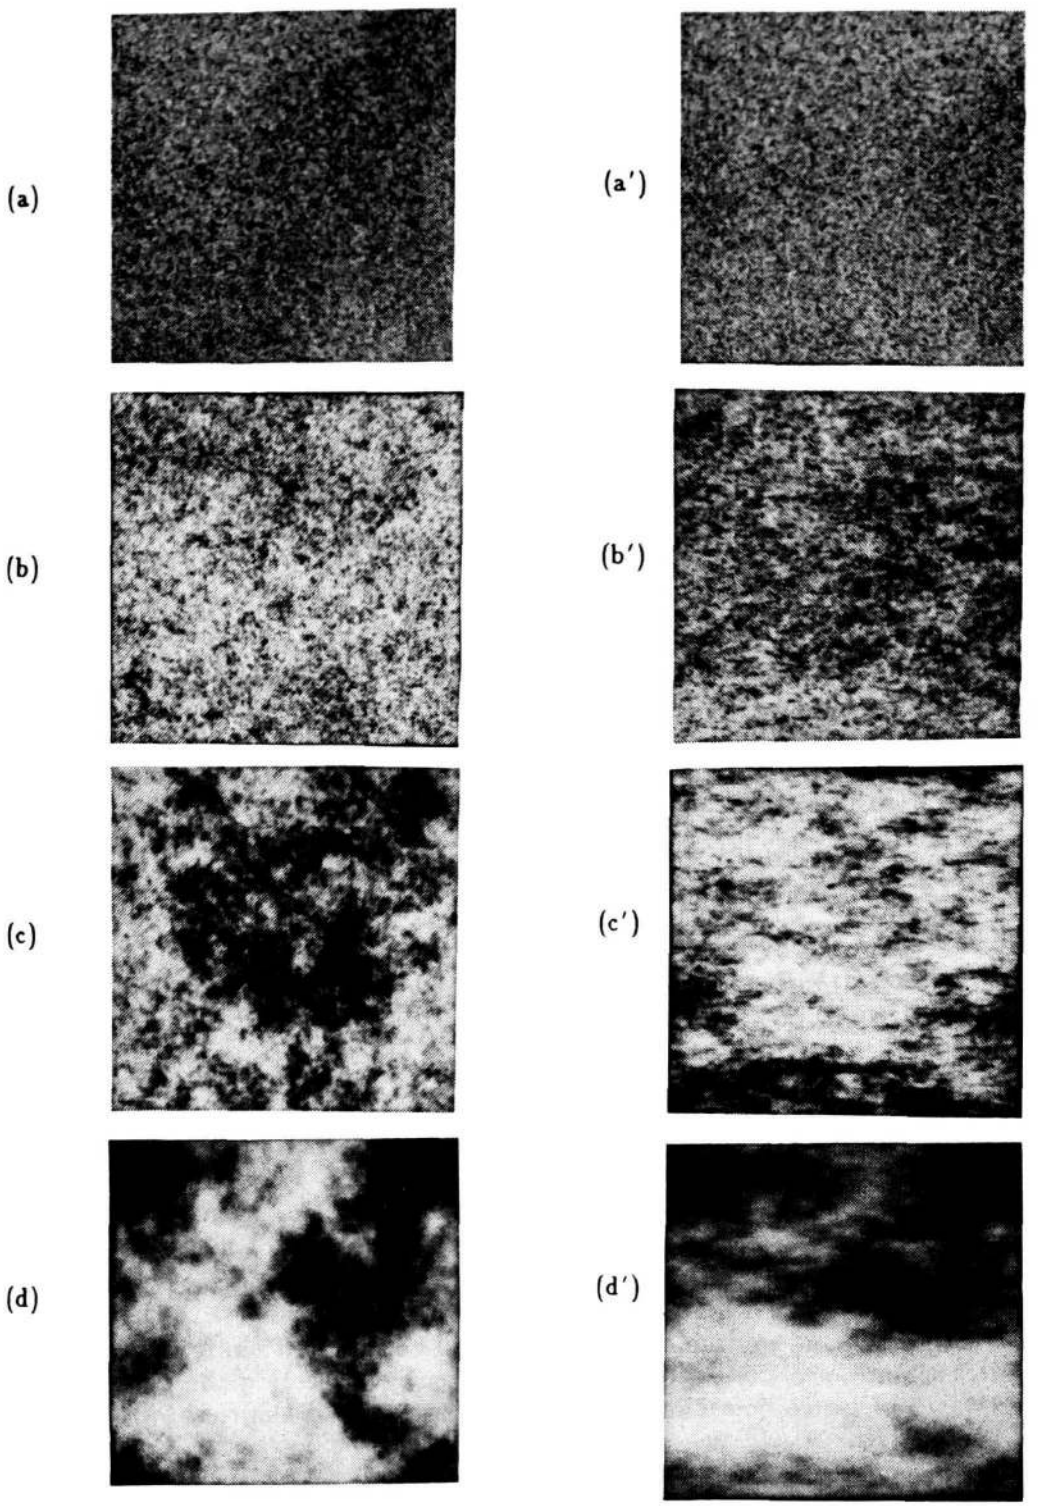

Figure 2

(a)-(d): Images of scaling noise surfaces, $\beta = 0, 1, 2$, and 3 respectively; slant $\sigma = 0$.
(a')-(d'): Images of scaling noise surfaces, $\beta = 0, 1, 2$, and 3 respectively; slant $\sigma = 70$ degrees.

scaling noise surfaces slanted at 70 degrees are shown in Figure 2(a'-d'); their anisotropy should be evident at least when β is large.) By the similarity theorem (cf. (1) above) the power spectrum of the image $I_{\beta,\sigma}$, which in the unslanted case $\sigma=0$ was

$$F_\beta(\omega,\theta) = k\ \omega^{-\beta},$$

becomes, for $\theta=\pi/2$ (i.e., parallel to the foreshortened y-axis):

$$F_{\beta,\sigma}(\omega,\pi/2) = \cos^2\sigma\ k\ (\omega\cos\sigma)^{-\beta},$$

and, for $\theta=0$ (i.e., normal to the foreshortened direction):

$$F_{\beta,\sigma}(\omega,0) = \cos^2\sigma\ k\ \omega^{-\beta}.$$

The ratio of the power in these two orientations, as a function of spatial frequency, is then[3]

$$R_{\beta,\sigma}(\omega) = \frac{F_{\beta,\sigma}(\omega,\pi/2)}{F_{\beta,\sigma}(\omega,0)} = (\cos\sigma)^{-\beta}, \tag{3}$$

and we can write

$$\sigma = \arccos\exp\left(\frac{\log R_{\beta,\sigma}(\omega)}{-\beta}\right). \tag{4}$$

Thus, we can determine the slant of the imaged surface if we can measure both β and $R_{\beta,\sigma}$. But this is easy to do, given the availability of frequency and orientation tuned channels of the sort introduced above. Say we have four channels with orientations at 0 and $\pi/2$ centered at each of two any distinct frequencies ω and $\alpha\omega$. Then, since α is known, a good estimate of β can be obtained from either the fact that

$$\frac{P_{\omega,0}(I_{\beta,\sigma})}{P_{\alpha\omega,0}(I_{\beta,\sigma})} \approx \frac{F_{\beta,\sigma}(\omega,0)}{F_{\beta,\sigma}(\alpha\omega,0)} = \frac{\cos^2\sigma\ k\ \omega^{-\beta}}{\cos^2\sigma\ k\ (\alpha\omega)^{-\beta}} = \alpha^\beta$$

or

$$\frac{P_{\omega,\pi/2}(I_{\beta,\sigma})}{P_{\alpha\omega,\pi/2}(I_{\beta,\sigma})} \approx \frac{F_{\beta,\sigma}(\omega,\pi/2)}{F_{\beta,\sigma}(\alpha\omega,\pi/2)} = \frac{\cos^2\sigma\ k\ (\cos^2\sigma\omega)^{-\beta}}{\cos^2\sigma\ k\ (\alpha\cos^2\sigma\omega)^{-\beta}} = \alpha^\beta\ ;$$

and then, with β known, and $R_{\beta,\sigma}$ estimated directly from

$$\frac{P_{\omega,\pi/2}(I_{\beta,\sigma})}{P_{\omega,0}(I_{\beta,\sigma})} \quad \text{or} \quad \frac{P_{\alpha\omega,\pi/2}(I_{\beta,\sigma})}{P_{\alpha\omega,0}(I_{\beta,\sigma})},$$

the slant σ can be determined by (4).[4]

Since the channels in the model have appreciable orientation bandwidth, they respond to image power in a range of orientations. This means that $P_{\omega,\pi/2}(I_{\beta,\sigma})$ will tend to be somewhat less than $F_{\beta,\sigma}(\omega,\pi/2)$, and $P_{\omega,0}(I_{\beta,\sigma})$ will tend to be somewhat larger than $F_{\beta,\sigma}(\omega,0)$, which in turn means that the estimate of $R_{\beta,\sigma}$ obtained in this way will be smaller than it should be, and so will lead to an underestimate of slant σ. But as can be seen by inspecting the results of simulating the model on images of isotropic scaling noise surfaces at various slants shown in Figure 3, the errors are not substantial.[5]

---

[3] The fact that $R_{\beta,\sigma}$ depends not on ω is the reason for calling these marking functions *scaling* noises.

[4] In practice, $P_{\omega,\theta}$ can be observed at a large assortment of frequencies, and the resulting estimates of β and σ can be subjected to averaging or a more sophisticated type of evidence composition.

[5] The case of the white noise surface is excluded; $\beta=0$ leaves (4) undefined. For such a surface, slanting does not break the isotropy of its image's power spectrum, and only changes its contrast: it is theoretically impossible to determine the slant of such a surface from a normalized image. Perhaps this is what Pentland ((Pentland 1984), p. 673) has in mind when he says "Characterization of an image in terms of radial slices of the Fourier domain ... constrains the shape of the 3-D surface hardly at all ... Illumination effects can account for most variation in

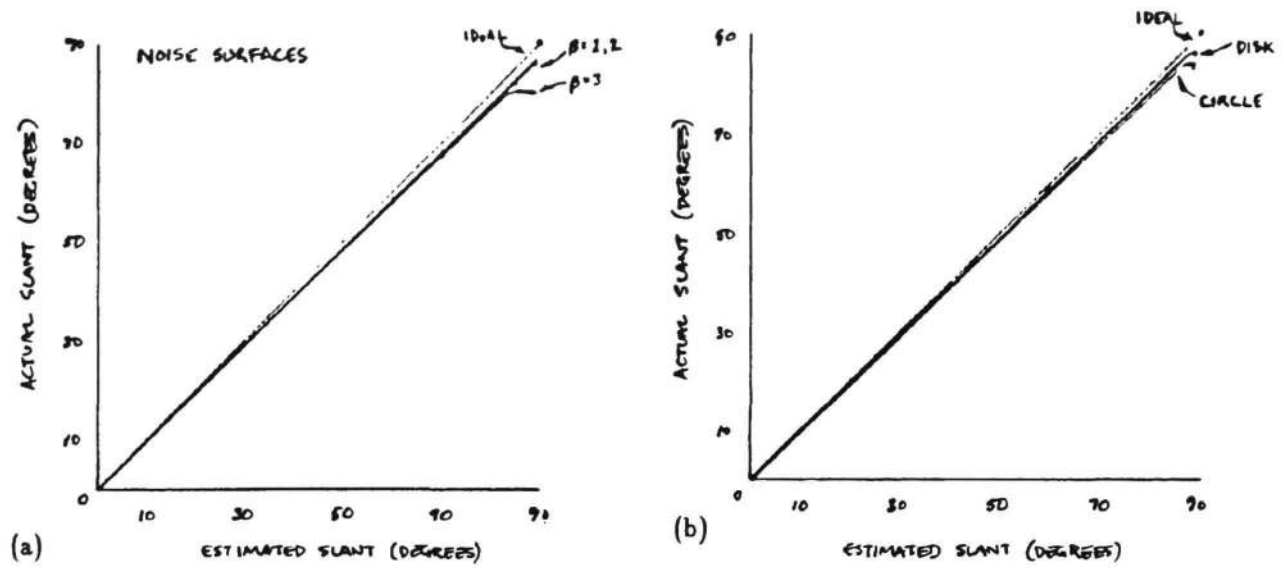

Figure 3

Estimates of slant generated by the 'adaptive β' model for (a) scaling noise surfaces, and (b) circles and disks.

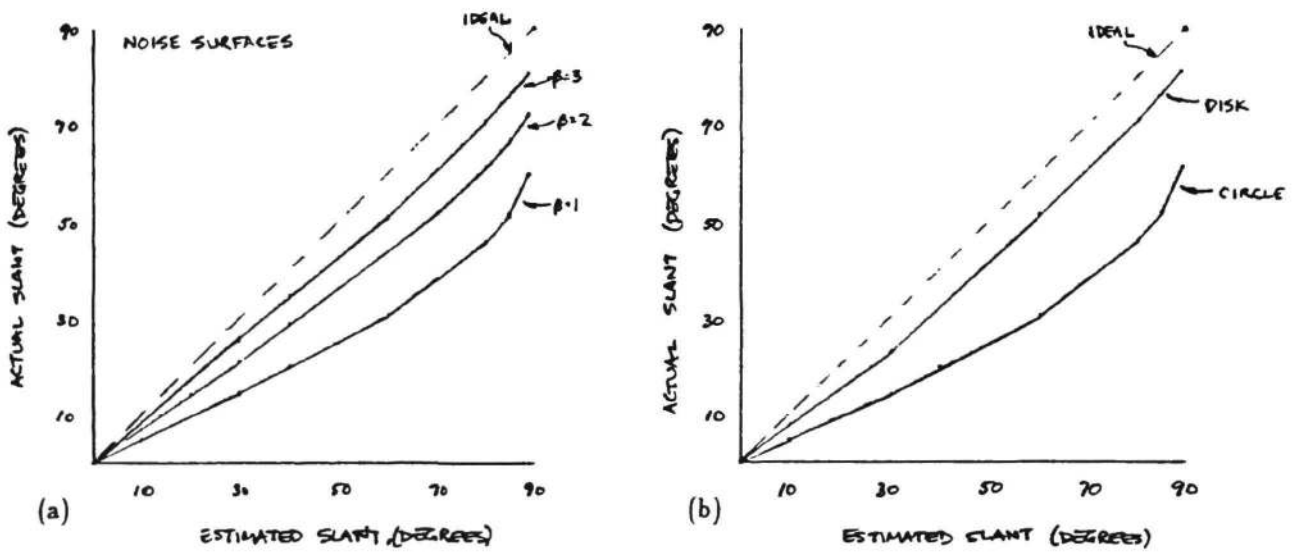

Figure 4

Estimates of slant generated by 'nonadaptive β' model based on Equation 5, $\beta_0 = 4$; for (a) scaling noise surfaces, and (b) circles and disks.

DETERMINING THE SLANT OF CIRCLES AND DISKS

The results of the previous section are not applicable only to stochastically textured surfaces; any surface marked with a function whose power spectrum is approximately proportional to $\omega^{-\beta}$ can have its slant estimated in the same way from a power spectrum of its image. In this section, we demonstrate how this works for isolated circles and disks.

Let a planar environmental surface be marked with a circle of radius r, ${}_C I_r(x',y') = \delta((x'^2+y'^2)^{\frac{1}{2}}-r)$, where δ is the unit impulse function. The orthographically projected image of this surface at zero slant is ${}_C I_r(x,y)$ which has power spectrum

$$F_r(\omega,\theta) = (2\pi r J_0(2\pi\omega r))^2 \quad .$$

Here J_0 is the Bessel function of the first kind of order zero, which can be closely approximated when $x > \pi$ by (cf. (Kreyszig 1972), p. 129)

$$J_0(x) \approx \frac{1}{\sqrt{\pi x}}\sqrt{2}\,\sin(x+\pi/4) \quad ;$$

so, if ω is not too small, we have

$$F_r(\omega,\theta) \approx \frac{4r}{\omega}\sin^2(2\pi\omega r+\frac{\pi}{4}).$$

Thus, an unslanted circle's power spectrum has an envelope like $k\omega^{-\beta}$ with $\beta=1$; and though the phase of the spectrum is not random as it is for scaling noise surfaces, the same tuned-channel techniques will work for slanted circles. The projection of a circle slanted at angle σ is the ellipse ${}_C I_r(x,y/\cos\sigma)$, which has power spectrum in the relevant directions $\theta=0,\pi/2$

$$F_{r,\sigma}(\omega,\pi/2) = \cos^2\sigma \quad \frac{4r}{\omega\cos\sigma} \quad \sin^2(2\pi\omega r\cos\sigma+\frac{\pi}{4})$$

$$F_{r,\sigma}(\omega,0) = \cos^2\sigma \quad \frac{4r}{\omega} \quad \sin^2(2\pi\omega r+\frac{\pi}{4}).$$

Now the ratio $F_{r,\sigma}(\omega,\pi/2)/F_{r,\sigma}(\omega,0)$ will be undefined for some ω since the denominator will vanish, but the averaging effect of the tuned channels (due to their having appreciable frequency bandwidth) will prevent the observation of zeros in $F_{r,\sigma}$ and will extract the envelope of the spectrum well enough to determine $\beta=1$ and then σ in the same way as for scaling noise surfaces.

A disk

$${}_D I_r(x,y) = \begin{cases} 1 & \text{when } x^2+y^2<r, \\ 0 & \text{otherwise} \end{cases}$$

has a $\omega^{-\beta}$ power spectrum envelope as does a circle of the same radius ${}_C I_r(x,y)$, but the scaling exponent β is, somewhat surprisingly, very different in the two cases. The disk has power spectrum

$$F_r(\omega,\theta) = (\frac{r}{\omega}J_1(2\pi\omega r))^2 \quad ,$$

where J_1 is the Bessel function of the first kind of order one, approximated when $x > \pi$ by

$$J_1(x) \approx \frac{1}{\sqrt{\pi x}}\sqrt{2}\,\sin(x-\pi/4) \quad .$$

Thus the power spectrum can be approximated by

$$F_r(\omega,\theta) \approx \frac{r}{\pi^2\omega^3}\sin^2(2\pi\omega r-\frac{\pi}{4})$$

---

such a description." But this is *not* true for values of β other than 0.

which has an envelope like $k\omega^{\beta}$ with $\beta=3$, compared to the circle's $\beta=1$. Again, the same tuned channel techniques permit an accurate estimation of the disk's slant. The results of a simulation of the model on these figures are displayed in Figure 3.

DISCUSSION OF THE MODEL

How suitable is this slant detection model as a model of human visual perception? One striking feature is its unrealistically accurate performance on images of scaling noise surfaces with small β: the model performs virtually flawlessly (cf. Figure 3) over nearly the entire range of slants when $\beta = 1$, but such images seem to produce very weak slant illusion for human observers and, even at a slant of 70 degrees do not have very salient anisotropies. As β increases, however, the salience of anisotropy seems to increase until, at $\beta=3$, Figure 2(d') looks fairly convincingly like a *slanted* patch of sun-dappled lawn.

This suggests that our model's 'adaptive' determination of the scaling factor β tailored for each image is unrealistic. Perhaps instead of (3), the visual system has

$$\frac{F_{\beta,\sigma}(\omega,\pi/2)}{F_{\beta,\sigma}(\omega,0)} = (\cos\sigma)^{\beta_0} \tag{5}$$

or

$$\frac{F_{\beta,\sigma}(\omega,\pi/2)}{F_{\beta,\sigma}(\omega,0)} = \frac{\beta_0}{\cos\sigma}$$

or

$$\frac{\pi}{2}\left(1-\frac{1}{1+\beta_0\,(F_{\beta,\sigma}(\omega,\pi/2)-F_{\beta,\sigma}(\omega,0))}\right) = \sigma$$

or some other monotonic mapping between a measure of anisotropy in the frequency plane and the possible range of slant angles with *fixed* parameter β_0. Given data from the psychophysical experiment, it would be possible to fit one or another such function; as an example, the results of a simulation of (5) with $\beta_0=4$ are shown in Figure 4, and they are closer than the adaptive model to what one might expect from human performance.

A problem with any such modification, however, is that to distinguish between the slant-detectability of scaling noise surfaces with $\beta=1$ and $\beta=3$ to fit human performance is *ipso facto* to distinguish to the same extent between the slant-detectability of circles and disks.[6] And while the experiments have not been done, it seems likely that circles and disks are very similar in respect of slant-detectability, and $\beta=1$ and $\beta=3$ noise surfaces are very different.

I suggest that the conclusion to draw is that the power in areas of an image's global frequency spectrum does not exhaust the information available to the perceptual system; spatial information also exists. In fact, as has been noted elsewhere,[7] if spatial frequency analysis is performed on the retinal image, it is performed locally, not globally; and the orientation- and

[6] This is true if the lowest spatial frequencies are ignored. At frequencies low enough to make the Bessel function approximations used here innacurate, or low enough to make the undulations in the circle and disk spectra observable by octave-bandwidth channels, a spatial frequency model could distinguish between noise on the one hand and circles and disks on the other. But low spatial frequencies are not the issue; low pass filtering the surfaces to remove these frequencies appears to change the slant-detectability of a large circle hardly at all, while practically eliminating the already scarcely detectable anisotropies in an image of a slanted $\beta=1$ noise surface. Anisotropies in images of a $\beta=3$ surface remain salient.

[7] For example, in (Daugman 1984a).

frequency-tuned channels that exist are almost optimally designed, from an information-theoretic point of view, to carry both spatial and frequency information. Now a scaling noise surface has a very uninteresting spatial organization; practically all its distinctiveness comes from the statistics of its radiance distribution, which is captured well in its power spectrum. Circles and disks and their projections, however, have a determinate shape, and it should play a role in the judging of slant. A parameterized tuned channel model of the sort sketched here may tell the story of slant detection for stochastic surfaces, but it will not be the whole story of slant detection.

Acknowledgements

The research reported here was supported by a grant from the Alfred P. Sloan Foundation to the Institute of Cognitive Studies at Berkeley. I want to thank Steve Palmer, Jitendra Malik, Gene Switkes and John Kruschke for helpful discussions.

REFERENCES

Bracewell, R. N. (1978). *The Fourier Transform and Its Applications*, McGraw Hill, 1978.

Campbell, F. W. & J. G. Robson (1968). Application of Fourier analysis to the visibility of gratings, *J. Physiol. 197* (1968), 551-566.

Campbell, F. W., G. F. Cooper, J. G. Robson & M. B. Sachs (1969). The spatial selectivity of visual cells of cat and the squirrel monkey, *J. Phsiol. 204* (1969), 120-121.

Daugman, J. G. (1984). Spatial Visual Channels in the Fourier Plane, *Vision Res. 24*, 9 (1984), 891-910.

Daugman, J. G. (1984). Representational Issues and Local Filter Models of Two-dimensional Spatial Visual Encoding, in *Models of the Visual Cortex*, D. Rose & V. G. Dobson (editor), Wiley, 1984.

DeValois, R. L. & K. K. DeValois (1980). Spatial Vision, *Ann. Rev. Psychol. 31* (1980), 309-341.

DeValois, R. L., E. W. Yund & N. Hepler (1982). The Orientation and Direction Selectivity of Cells in Macaque Visual Cortex, *Vision Res. 22* (1982), 531-544.

Ginsburg, A. P. (1978). *Visual information processing based on spatial filters constrained by biological data*, Cambridge University, 1978. Ph.D. dissertation.

Janez, L. (1983). Stimulus control of visual reference frame orientation: Quantitative theory, *Informes de Psicologia*, 1983, 133-147.

Kreyszig, E. (1972). *Advanced Engineering Mathematics*, John Wiley and Sons, 1972.

Maffei, L. & A. Fiorentini (1973). The visual cortex as a spatial frequency analyser., *Vision Res. 13* (1973), 1255-1267.

Mandelbrot, B. B. (1983). *The Fractal Geometry of Nature*, W. H. Freeman, New York, 1983.

Pentland, A. P. (1984). Fractal-based description of Natural Scenes, *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-6*, 6 (1984), 661-674.

Sachs, M. B., J. Nachmias & J. G. Robson (1971). Spatial frequency channels in human vision, *J. Opt. Soc. Am. 61* (1971), 1176-1186.

Schiller, P. H., B. L. Finlay & S. F. Volman (1976). Quantitative studies of single-cell properties of monhkey striate cortex , *J. Neurophysiol. 39* (1976), 1288-1351.

Watson, A. B. (1983). Detection and Recognition of Simple Spatial Forms, in *Physical and Biological Processing of Images*, O. J. Braddick & A. C. Slade (editor), Springer Verlag, Berlin, 1983.

Witkin, A. P. (1981). Recovering Surface Shape and Orientation from Texture, *Artificial Intelligence 17* (1981), 17-45.