Chapter 3

Image Formation

3.1 Introduction

The goal of this thesis is the development of a classifier which discriminates between rough surfaces on the basis of their visual appearance. This chapter links the surface description developed in the last chapter to the image incident on the camera lens. It will be shown that the characteristics of this link depend on, among other things, the illuminant tilt. This chapter develops a model of the source/surface dependency to which we aim to make the classifier robust.

The aim of this chapter is to model the transition from physical surface to incident image. This aim is attained by satisfying two objectives: firstly, the formation of a description of the mapping from surface derivatives to image intensity; and secondly, the development and evaluation of a frequency based transfer function from surface to image. Whereas the first is a purely local process, the second takes into account the spectral characteristics of the surface and forms the link between the spectral surface models described in the previous chapter, and the image, data set and frequency-based feature set considered in subsequent chapters.

The first aspect of the model is considered in section 3.2; a brief description of the underlying physical mechanisms of reflection and some of the more prominent developments in reflectance modelling, is given. The thesis then focuses on the much less researched area of diffuse reflection: three modern models are considered, two of these, as well as the classical Lambertian model, are compared with an empirical mapping. We will find the analytical function which most closely resembles the measured function under the conditions associated with rough surfaces.

The second aspect of the model is a frequency-based transfer function between surface and image—again taking an analytical model as the starting point. In this section we ask two questions:

- is the analytical model optimal in the least squares sense?
- how does the nature of the surface affect the accuracy of estimation?

This section therefore considers the optimality of the analytical model as well as the accuracy of a more empirical linear model.

Finally, the predictions of the analytical model are compared with the actual characteristics of an isotropic and an anisotropic texture.

3.2 Modelling the Reflectance Function

In this section we shall consider the purely local interaction between a surface facet and lighting geometry to give the observed intensity. This problem has been investigated by researchers in physics [Torrance67], computer graphics [Cook82] and more recently machine vision [Nayar91]. The phenomenon of reflection can be produced by one or more different physical processes. These processes give rise to various generic behaviours, as well as intermediate forms. In this section we will briefly discuss the different characteristic behaviours, as well as the physical processes behind them, before focusing on one generic type. Several models of this type will be evaluated against the reflection characteristic of the surfaces used in the experiments described in this thesis. The findings of this section will then be used in the next to produce a non-local model of illumination, which will link the surface models of the previous chapters to the feature models of Chapter 5.

3.2.1 Terms used in this chapter

We begin by stating the standard radiometric definitions. These are given in more complete terms in most textbooks [Watt p.91] and many papers [Nayar91], we reiterate them here for completeness. *Flux* is the rate of emission or reflection of light energy, and is measured in Watts. *Irradiance* is the incident flux per unit surface area (Wm^{-2}), *radiant intensity* is the flux radiated per solid angle in a particular direction ($Wm^{-2}\ St^{-1}$) . *Radiance* is the radiant intensity, per unit, of *foreshortened* area ($Wm^{-2}\ St^{-1}$), where the foreshortened area is the surface area times the cosine of the angle between the radiated light and the surface normal. For a Lambertian surface the radiance is independent of the viewing angle. Horn has shown that image irradiance, i.e. the pixel intensity value is proportional to the scene radiance [Horn86].

The ratio of surface radiance to incident irradiance is defined as the bi-directional reflection distribution function (BRDF) (equation 3.2.1a), which Horn attributes to Nicodemus. For a facet with isotropic microstructure, the BRDF is invariant to rotation of the surface about the surface normal and may be expressed in the form shown in (3.2.1b) [Horn p.210].

$$f(\theta_r,\theta_i,\eta_r,\eta_i) \tag{3.2.1a}$$

$$f(\theta_r,\theta_i,\eta_r\text{-}\eta_i) \tag{3.2.1b}$$

where θ is the tilt angle and η is the slant angle. The subscript denotes whether the angle refers to the incident (i) or reflected (r) ray.

In this thesis, we will borrow the concept of the reflectance map, developed by Horn [Horn77] and widely used in *shape from shading* research. The reflectance map assumes a specific lighting orientation and describes the radiance/irradiance ratio in terms of the surface normal. It is therefore a two dimensional function with two independent variables, the directional derivatives. The measured intensity is the dependent variable.

While the reflectance map is stated in terms suited to our purposes, the fact that it has two independent variables makes it difficult to show graphically and we seek to describe the effect with a single variable. The Lambertian model has the assumption that reflectance is a function purely of the angle of incidence and is independent of viewing direction. This assumption has been challenged in both [Oren94] and [Wolff94], though Healey's model does retain the assumption [Healey89]. We will show the reflectance function as a function of incident angle averaged over the relevant range of viewing angle. That is, the reflectance function is calculated in terms of surface gradient and expressed in terms of incident angle.

3.2.2 Underlying Physical Processes

Light incident on a flat surface will be reflected and transmitted. The degree to which either occurs depends on the electrical characteristics of the material. In a conductor electrons are loosely bound to atoms and will oscillate at the same frequency as the incident field. The field itself will be quickly attenuated as it penetrates the material, however, the surface currents induced by the field will in turn reradiate an electromagnetic field in the mirror direction. If on the other hand, the electrons are tightly bound to atoms, as they are in a dielectric, they will interact with the field to a much lesser degree and the light will, in the main, be transmitted. We would therefore expect reflection to be the dominant effect in conductors and transmission in dielectrics.

In fact, dielectrics are rarely transparent; light *is* reflected from the material, though by a different mechanism and in a much less directional manner than with the surface reflection associated with conductors. The actual mechanism is unknown, though the most popular hypothesis is that of ‘bulk (or body) scattering’. This model assumes the material to be inhomogeneous; light incident on the surface is refracted at the boundary and is transmitted through the material until it encounters an inhomogenity boundary from which it is scattered. Light is repeatedly reflected and refracted from these internal boundaries and some is scattered in the direction of the air boundary and will emerge in a randomised direction. In the limiting ‘Lambertian’ case, the surface will appear equally bright from all viewer orientations.

The previous discussion attributed optical effects solely to the material characteristics and yielded the two extreme cases: specular reflection for conductors and Lambertian reflection for dielectrics. In fact, both conductors and dielectrics exhibit transitional behaviour. We consider surface reflection from dielectrics first. The ratio of surface reflected light to incident light is given for both dielectrics and conductors in terms of the Fresnel equations. These are defined in terms of the complex refractive index, $n+iK$, and the angle of incident light with respect to the surface normal. If we plot the Fresnel expression for unpolarised light incident on a smooth dielectric, we obtain the curve shown in *Figure 3.2.1*. The surface reflection rises sharply as the incident light approaches the grazing angle and, assuming a conservative system, the total body reflection will decrease accordingly.

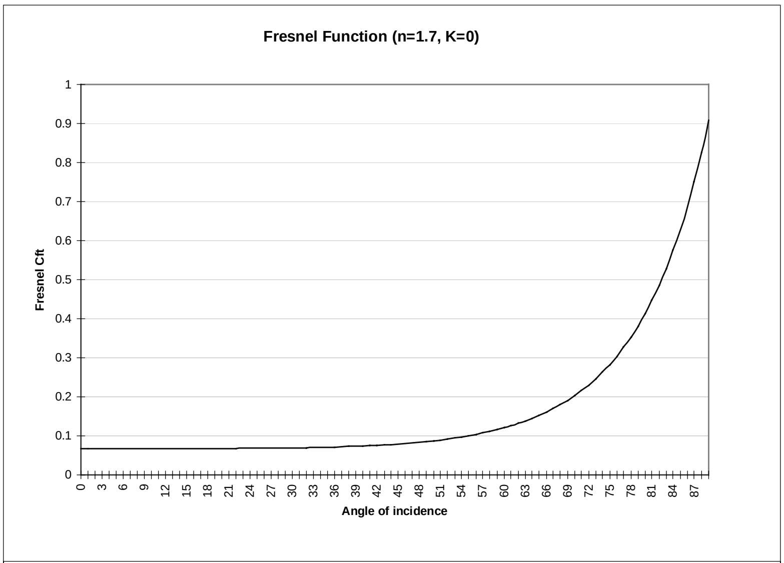

Figure 3.2.1 Fresnel Function

Just as dielectric reflection shows some characteristics associated with conductors, conductor reflection may also fall into the transitional region. Although metals attenuate transmitted light far too quickly for body reflection to be significant, surface roughness will cause the reflected light to be more diffuse. A rough surface no longer has a single mirror direction, but instead a statistical distribution of directions, that manifests itself as a variation of intensity. This effect will be most pronounced in the conductor case, though it will also have an effect for dielectric reflection.[a1]

3.2.3 Specular Models

The most general methodology for dealing with surface reflection, both in terms of the degree of roughness and surface conductivity, is scattering theory. A less general and rigorous technique, which models the surface as a collection of perfectly mirror-like microfacets whose slopes conform to a specified distribution was developed by Torrance and Sparrow [Torrance66]. The Torrance-Sparrow (TS) model is only relevant where the level of roughness is much greater than the wavelength of light. Although this restriction excludes surfaces with optical finish, the TS model forms the basis for the majority of specular models recently developed in both computer graphics and machine vision. We argue that the substitution of the scattering model by a TS type model is particularly valid in texture analysis given the large body of experimental evidence that the roughness spectrum is continuous to almost atomic scales.

The TS model consists of three sub-components: a slope distribution (analytically modelled as a distribution of V-grooves), a geometric factor used to model facet shading and a Fresnel coefficient. Blinn is the first to have applied the model to computer graphics, appending a Lambertian component [Blinn77]. Cook further developed the model by introducing a new slope distribution, defined in terms of real measurable parameters and introducing a spectral dependency [Cook82].

3.2.4 Diffuse Models

While specular reflection has been thoroughly investigated in the field of computer graphics, researchers in this area seem to be content to assume diffuse reflection to be adequately modelled by the Lambertian model. However, diffuse reflection is also widely encountered in machine vision and recently the need for accurate reflectance maps has motivated the investigation of diffuse reflection. We consider three analytical models developed in recent years, the model developed in [Oren94] is the least physically-based and, like the TS model it models the surface as a group of microfacets, although in Oren's model the facets are perfectly Lambertian, rather than specular, reflectors. Healey [Healey89] adapts existing theory of diffuse reflection and integrates it closely with the TS model of specular reflection. Wolff [Wolff94a] develops a more complex model, which includes a viewer orientation dependency. This is valid for smooth surfaces only, though in a later paper the model is extended to include rough surfaces. We will briefly discuss these models with reference to each other and the Lambertian model in the context of textured surfaces.

Lambertian Model

In intuitive terms, the Lambertian model states that the perceived intensity of a surface facet is dependent only on the relative geometry of that facet and an illuminant. Mathematically, this is represented as the dot product of the surface derivative vector with the illuminant vector.

While the radiance of the facet is constant with respect to the viewer's position, we note that the radiant intensity, that is the flux per unit solid angle, varies with the foreshortened area of the facet.

The radiance is plotted against the angle of incidence in *Figure 3.2.2*; as with all the other reflectance functions in this thesis the function is scaled to give unity radiance for a flat surface, i.e. at an angle of incidence of 60° where the slant angle is held constant at 60°. This is compatible with our experimental practice of using the image of a flat surface with a uniform albedo equal to that of the textures to remove variations in the irradiance incident on the surface.

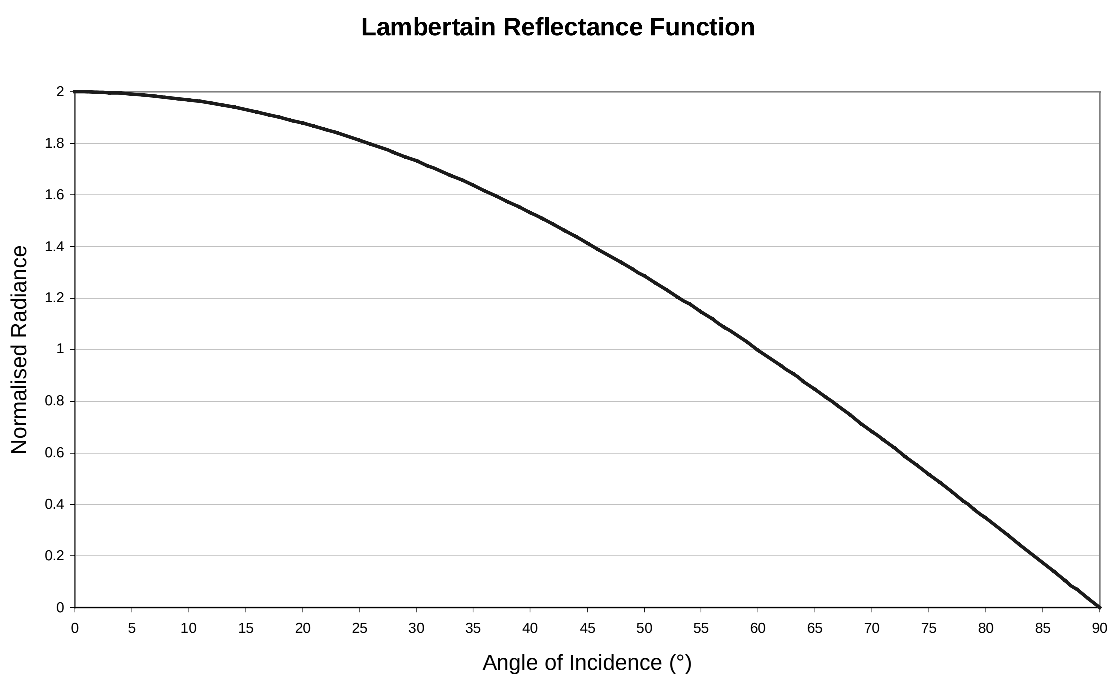

Figure 3.2.2 The Lambertian Reflectance Function

Oren's Reflectance Function

Oren states that the Lambertian model is inadequate for many applications and that this failure is most apparent when viewer direction (relative to the surface normal) is varied. It is Oren's thesis that the discrepancies are primarily due to subpixel surface roughness and he does not pursue the physics of diffuse reflection. Like the TS model, Oren models the surface as a set of long V-grooves consisting of two facets, although in Oren's model these are perfect Lambertian, rather than specular reflectors. Another difference is in the models' facet distributions: the TS facets are of equal area, and the distribution represents the number of facets with a specific normal that lie within a unit area. Oren's distribution, in contrast, represents the fraction of the area occupied by facets of a given normal; this, he argues, makes the model less sensitive to variations in the actual roughness distribution. Oren uses these facets to model shadowing, masking as well as interreflection, though this last term is subsequently neglected. Finally Oren develops a functional approximation for the model.

We plot Oren's reflectance function of radiance against angle of incidence for three degrees of sub-facet roughness *Figure 3.2.3*. This function, as with all the functions reported in this section, has been normalised at $\theta_i = 60°$, i.e. it assumes a value of unity for the radiance of light reflected from a flat surface

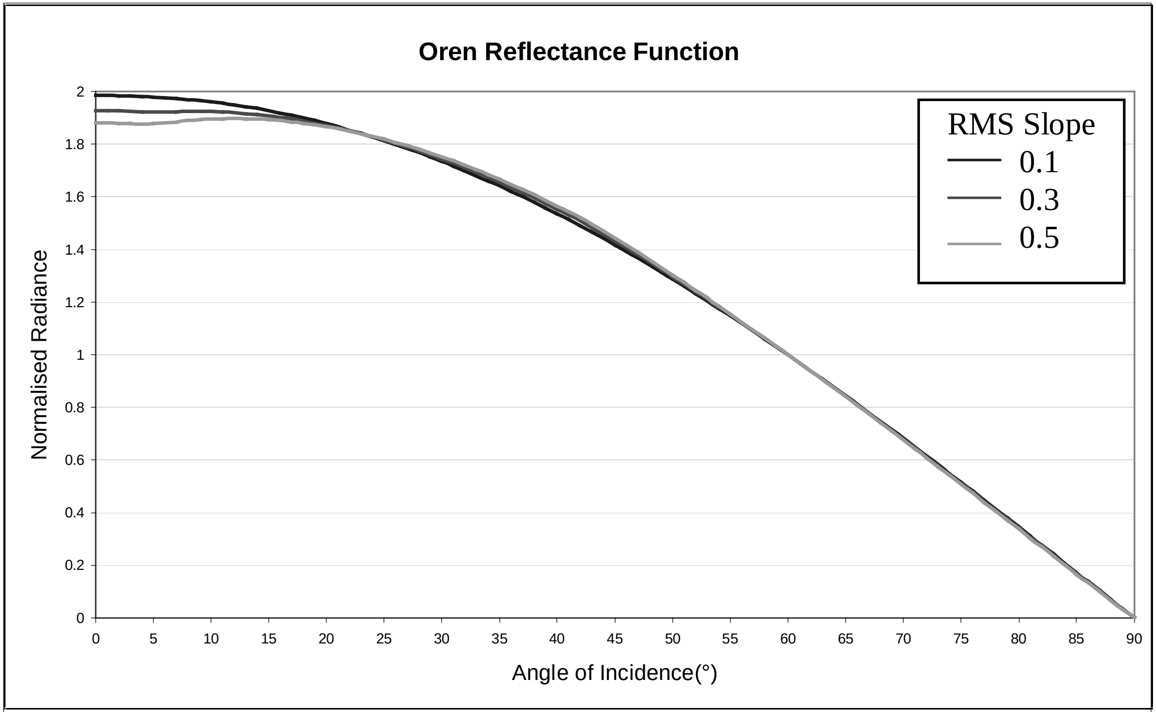

Figure 3.2.3 Oren's reflectance function for various degrees of surface roughness

Oren then experimentally verifies his model on a set of 2X2 inch samples imaged with a CCD camera at 6ft, with the average pixel value as the sample irradiance. He finds the model to be in strong agreement with his experiments for samples of wallplaster, painted sandpaper and white sand (though a small specular component was detected in this case). Other notable results are that as roughness increased, the reflectance map varied from the Lambertian to a more linear form. Another interesting result is that backscatter is greater than foreground irradiance—not only illustrating an inadequacy in the Lambertian model but also contrasting with the case of specular reflection from a rough surface. For samples of cloth, foam and woodshaving, Oren appends a TS model for specular reflection, with varying degrees of success. Several parameters of the model are estimated by fitting the data to the model using non-linear optimisation. He assumes the Fresnel coefficient to be unity due to difficulties in measuring *n*.

Healey's Reflection Model

Healey's dielectric model must be set in the context of the rest of his paper [Healey89], which develops a dichromatic reflection model with the aim of discriminating between materials based on their reflectance properties. He therefore considers metals as well as dielectrics and wavelength forms an important part of his model. This thesis is concerned only with dielectric surfaces and only considers monochrome images, wavelength dependencies are therefore neglected in the interest of clarity.

Healey models surface reflection using the TS model, with the Cook and Torrance slope distribution [Cook82]. He notes that for perfect dielectrics the refractive index is purely real, simplifying the Fresnel equations. Internally, scattering is based on a modified form of the Kubelka-Munk (KM) theory. The original KM theory treats the dielectric as consisting of many elementary layers, in which colorant particles are embedded. These layers are characterised by two parameters: α_h the fraction of light absorbed per unit path length and σ_h the fraction of light scattered per unit path length. The reflectance function predicted by the KM model is:

$$R_\infty = \frac{2 - w_h - 2[1 - w_h]^{0.5}}{w_h}$$

where

$$w_h = \frac{\sigma_h}{\sigma_h + \alpha_h}$$

In the case of Lambertian scattering, *w* is equal to unity.

The KM theory was extended by Reichmann [Healey89] to allow non-diffuse light, and vehicles with refractive indices which differ from those of air to be modelled:

$$R_b = (1 - R_s)\frac{C(\theta)(1 - r_i)[R_\infty - D(\theta)]}{2[1 - r_i R_\infty]\cos\theta}$$

where r_i is the internal diffuse surface reflectance, described by:

$$r_i = \frac{1 - 1.439 - 0.7099n + 0.3319n^2}{n^2}$$

where *n* is the real part of the refractive index,

and

$$C(\theta) = \frac{w_h \cos\theta(2\cos\theta + 1)}{1 - 4[1 - w_h]\cos^2\theta}$$

$$D(\theta) = \frac{2\cos\theta - 1}{2\cos\theta + 1}$$

It is worth reiterating that, in contrast to the models of Oren and Wolff, Healey's model does not incorporate any viewer dependency.

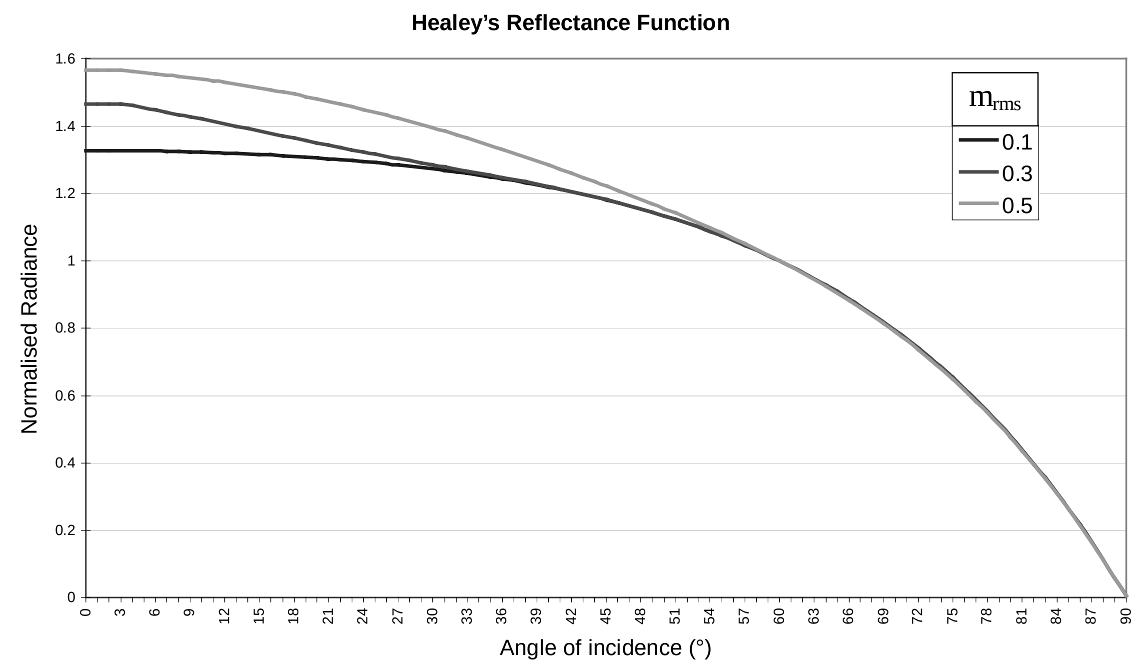

Figure 3.2.4 Healey's reflectance function for various surface roughnesses

Wolff's Reflectance Model

Wolff [Wolff94] also develops a model for sub-surface scattering, and in [Wolff94a] he extends this model to include rough surfaces and a surface reflection term by linearly adding the TS model. Wolff bases his model on the theory of Chandrasekhar; like the KM model this assumes a vehicle with refractive index equal to that of air. Wolff couples this model to the Fresnel boundary conditions, noting changes in the direction and subtended solid angle of the incident light.

Unlike the other models considered in this section, the Wolff model was not implemented due to its complexity. We are therefore unable to evaluate its effectiveness.

Comparison of Models

The relative characteristics of the illumination models, can be examined by rendering three spheres with Lambert's, Oren's and Healey's reflectance functions (see *Figure 3.2.5*). The most obvious difference between the renderings is in the regions of maximum reflectance. Lambert's function gives a distinct peak in intensity while Healey's model does not produce any discernible bright spot. Oren's model is intermediate between these forms.

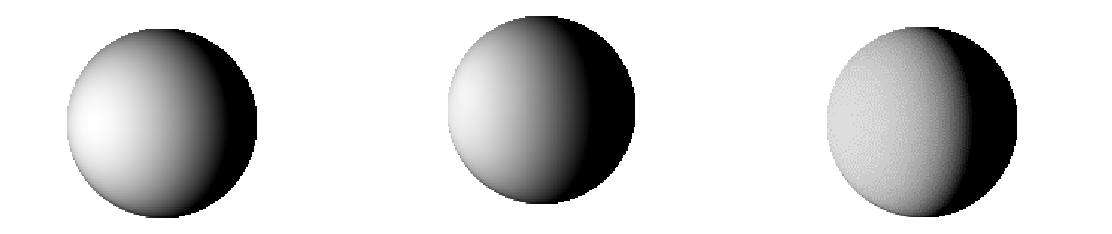

Figure 3.2.5 Spheres rendered with Lambertian (a), Oren(b) and Healey Models (c)

The highlight effect is also visible in the reflectance maps of the models (*Figure 3.2.6*). The lessening of the effect suggests that the relationship between surface derivative and radiance is more linear for Oren's model than Lambert's, as was noted by Oren himself.

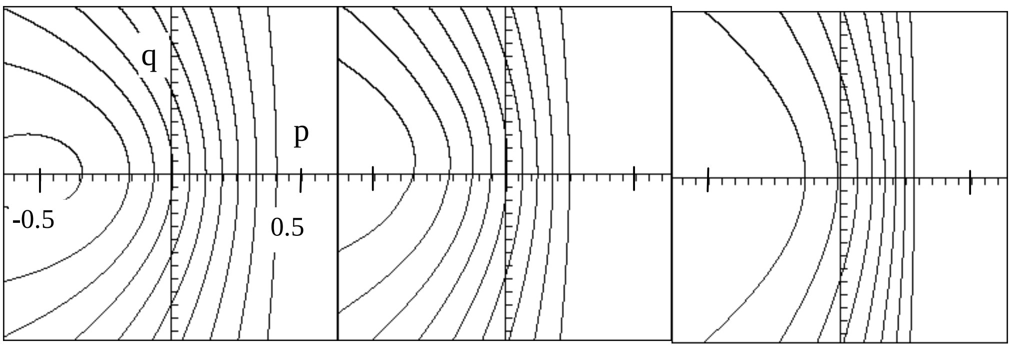

Figure 3.2.6 Reflectance Maps for (a)Lambert's, (b)Orens's and (c)Healey's Models

Comparison of the models with the empirically measured function is of more relevance to this thesis. All the surfaces considered in this thesis have been sprayed matt white. The reflectance map associated with this surface characteristic was derived by measuring the intensity of an object of known geometry, i.e. a sphere, and recording the mapping between surface derivatives and radiance. It is shown together with the analytical reflectance models in *Figure 3.2.7*.

It can be seen that Healey's model performs relatively poorly for the surface type used in this report. The Lambertian model is a much better approximation, however, it is Oren's model that is closest to the measured reflectance function.

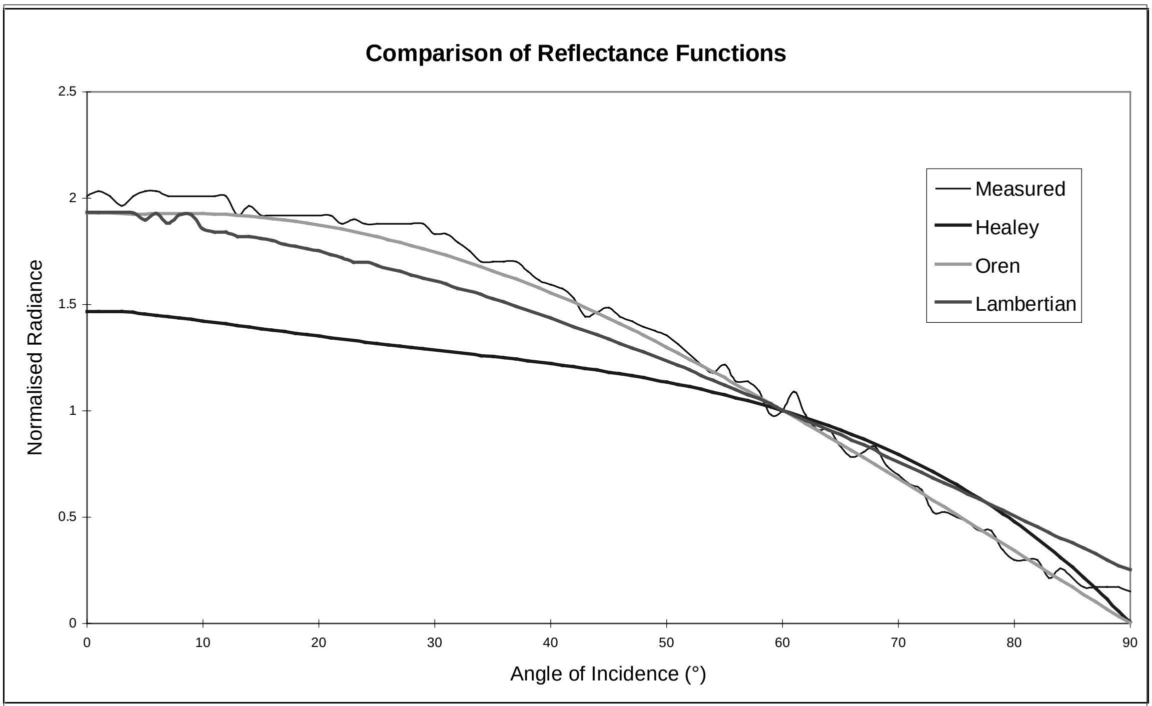

Figure 3.2.7 Comparison of Reflectance Function

While the comparison of reflectance functions is useful from an analytical point of view, it is also important to assess how well the reflectance functions perform on rough surfaces of the type considered in this thesis. The accuracy of the reflectance function will vary with incident angle. Light incident on a rough surface will have a distribution of incident angles, which will be partially dependent on the slope distribution. In order to assess the accuracy of the models for textured surfaces we must take the angular distribution into account. We do so by applying the models to synthetic surfaces with different slope distributions. This approach will therefore give a better indication of the actual accuracy of the function for our application. We evaluate the accuracy by rendering a series of synthetic height maps using an experimentally measured reflectance map. These results are then compared with analytical estimates derived using the three models. A signal to residue ratio (S/R) is used to perform this comparison. It is defined as

$$S / R = 10 Log_{10}\left(\frac{Var[i(x,y)]}{Var[e(x,y)]}\right)$$

where Var[*i(x,y)*] is the variance of the image predicted using the empirical function and Var[$e(x,y)$] is the variance of the difference between this, and the image predicted by the relevant theoretical model.

The first group of surfaces we consider is based on an isotropic fractal surface which is scaled to give varying degrees of roughness. The results show the Lambertian function to be the most accurate model for the reflectance function of our experimental surfaces, followed closely by Oren's model, with Healey's model performing comparatively poorly on our surfaces (see *Table 3.2.1)*

	S/R (dB)		
RMS Slope	**Lambertian**	**Oren**	**Healey**
0.125	18.96	17.05	12.12
0.250	18.77	17.95	10.27
0.500	16.50	15.81	12.10

Table 3.2.1 The accuracy of image prediction from isotropic surfaces of various roughnesses.

We now consider the highly directional Ogilvy surface defined in the previous chapter. In the first case, the grain of the surface runs at right angles to the direction of illumination, i.e. the illuminant is aligned with the direction of maximum slope variance *Table 3.2.2*. We note an increase in accuracy over the isotropic case, however the fall in accuracy associated with increasing roughness is still apparent.

RMS Slope		**S/R (dB)**		
$\mathbf{p_{rms}}$	$\mathbf{q_{rms}}$	**Lambertian**	**Oren**	**Healey**
0.125	0.034	20.43	18.76	12.42
0.250	0.068	20.62	21.30	10.01
0.500	0.136	16.26	16.26	12.39

Table 3.2.2 The accuracy of image prediction for directional surfaces, illuminant perpendicular to surface grain.

If we illuminate the same series of directional surfaces in the same direction as the surface grain, we find a significant fall in the accuracy of the Lambertian, and even more so in the Oren functions, whereas Healey's model is relatively unaffected (*Table 3.2.3*) . The poor performance of the Lambertian and Oren models is not solely an effect of the low slope in the direction of illuminant. We infer this from the fact that the maximum *rms*. slope in the direction of the illuminant for this category of surface (p_{rms} =0.136) exceeds the minimum slope parallel to the illuminant for two previous surface types (p_{rms} =0.125), yet the accuracy is still substantially below that of the first two cases.

RMS Slope		**S/R (dB)**		
p_{rms}	q_{rms}	**Lambertian**	**Oren**	**Healey**
0.034	0.125	11.13	2.18	11.12
0.068	0.250	13.66	9.87	12.48
0.136	0.500	14.22	13.47	10.82

Table 3.2.3 Accuracy of prediction for directional surface illuminated in the direction of surface grain.

We conclude that of the models tested here, the reflectance function of the surfaces used in this thesis is best approximated using the Lambertian model. The model performs well for isotropic surfaces and directional surfaces of low slope, which have their axis of maximum slope variance aligned with the illuminant. The Lambertian model performs most poorly on directional surfaces of low slope when the illuminant is parallel to the grain of the surface. For the directional surfaces considered in this thesis, the accuracy of the Lambertian function never falls below 11dB while for isotropic surfaces it remains above 16.5 dB.

3.2.5 Summary

In this section we have reviewed three analytical models of diffuse reflection, and evaluated these against the empirically measured reflectance map for our test surfaces. We found that, as a function of incident angle, Oren's was the closest to our measured function. If the test is weighted by the frequency with which facet orientations are predicted by our surface models, however, it was found that the Lambertian function modelled the empirical mapping more accurately. We found that the Lambertian model performed best on isotropic surfaces with low slope angles and on directional surfaces of low slope when the illuminant is perpendicular to the surface grain. The model performed most poorly on directional surfaces of low slope when the illuminant is parallel to the surface grain. However, even in the worst case, the S/R ratio of the Lambertian model never fell below 11 dB for the range of surfaces considered in this report.

3.3 An Optimal Linear Model of Image Formation

In the previous section we treated the reflectance function as a purely local phenomenon. However, any work concerned with texture is inherently concerned with the interaction of a point with its neighbours. In the previous chapter we modelled this interaction in the frequency domain. The radiance of a facet is a function of its surface derivatives—the spectra of which are easily related to the surface height spectrum. It is therefore attractive to model the transition from surface to image in the frequency domain.

For a tractable frequency domain analysis we require a linear model of the rendering process. If we assume a single reflectance function is valid over the entire surface, and that the illumination conditions are constrained, we can state that the only factor that will affect the degree to which the transfer function can be described as being linear will be the topographical characteristics of the surface. These will affect:

- the area and location of the region of the reflectance map which the model must approximate, and
- the degree of cast and self-shadowing[1].

From this it follows that a linear approximation is more valid for some surfaces than others.

The first objective of this section is to identify the degree of accuracy with which a linear model can represent the rendering of various surface types. This will be carried out in section 3.3.1 by using a least squares filter to obtain the optimal linear reflectance function for a given surface and illumination conditions. The surfaces will be scaled versions of two prototypes: the isotropic fractal, and the directional Ogilvy surface. The synthetic surfaces will be rendered using the empirical reflectance map and the least squares derived linear model. The criterion of accuracy will be stated in terms of the signal to residue ratio. This section will conclude with the statement of a number of constraints on the characteristics of surfaces which can be rendered, using an optimal linear model, to a specified S/R ratio.

---

[1] Shadowing may take two forms. A *cast* shadow occurs where one part of the surface prevents another part from being illuminated, by blocking the direct path between the light source and the shadowed area. *Self-shadowing* occurs when a facet is oriented such that it does not present an area on which light is incident, i.e. $\underline{\mathbf{I}}.\underline{\mathbf{N}}<0$. While the latter is a function of facet orientation and is a purely local phenomenon, the former is a function of surface height and is non-local in effect.

3.3.1 A Linear Filter

In this section we treat the problem as one of developing an optimum filter for a given illumination condition and also for a given surface. The optimum filter represents the transfer function from the surface derivatives to the image. Applying the filter to the surface derivative fields will give an approximation to the image. The filter will have the general form shown below:

$$i(x,y) = \mathbf{S}(x,y)^T \mathbf{V}$$

where:

$$\mathbf{S} = \begin{bmatrix} p(x,y) \\ q(x,y) \\ 1 \end{bmatrix} \qquad \mathbf{V} = \begin{bmatrix} a \\ b \\ c \end{bmatrix}$$

i(x,y) is the image intensity at point *(x,y)*,

p(x,y) and *q(x,y)* are the surface derivatives at *(x,y)*,

and *a,b* and *c* are the filter weights.

Let

$$\mathbf{R} = E[\mathbf{S}\mathbf{S}^T] \qquad \mathbf{U} = E[i_d \mathbf{S}]$$

where i_d is the desired signal.

To find the optimum parameter values we must minimise the quantity *e*, where

$$e = E[i_d^2] + \mathbf{V}^T\mathbf{R}\mathbf{V} - 2\mathbf{U}^T\mathbf{V}$$

The transfer function is approximated by a least squares linear filter in *p* and *q* calculated for the surface and image to which it will be applied.

3.3.2 The Signal to Residue Ratio

The analytical work of Chapter 5 is based on Kube and Pentland's frequency domain model [Kube88]. In order to link the surface and image spectra, Kube was obliged to use a linear approximation to the reflectance map. In this section we aim to assess the conditions under which the assumption of linearity is reasonable by comparing renderings using the optimal linear model *for a given surface and illuminant* with those of the empirically derived reflectance map. We also seek to discover the nature of the errors in spectral terms.

We will consider the transition from surface *s(x,y)* to images *i(x,y)* and *i'(x,y)* as being performed by process *o(p,q)*(derived experimentally using the matte-white sphere) and its linearised version *o'(p,q)* respectively. The difference between the desired signal *i(x,y)* and the linear estimate *i'(x,y)* is denoted by *e(x,y)*, the residual signal (*Figure 3.3.1*).

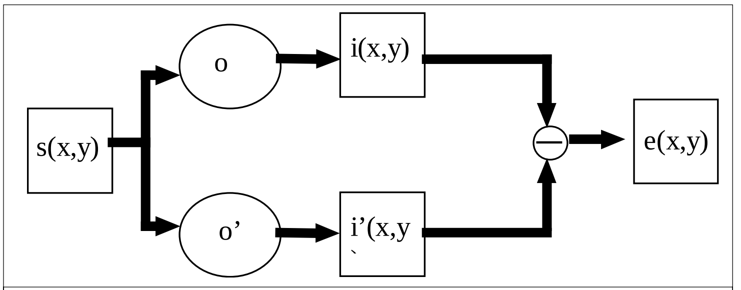

Figure 3.3.1 Calculation of residue image.

The signal to residue ratio is defined as:

$$S / R = 10 Log_{10}\left(\frac{Var[i(x,y)]}{Var[e(x,y)]}\right)$$

where Var represents the variance of a signal.

The surface vector field, ***S(x,y)***, is estimated from the potential field, *s(x,y)*, using two simple two-point estimators. The two point estimator is the simplest form of differentiator, and we note that it underestimates the magnitude of high frequencies, nevertheless we use the two-point estimator due to its simplicity, see [Bentum96] for a readable treatment of the subject.

3.3.3 The Effect of Surface Characteristics

Surface Roughness

We now consider the effect of surface roughness, parameterised by *rms slope*, on the accuracy of the linear model. An isotropic fractal surface was scaled and rendered using the empirical reflectance map. The accuracy of prediction is plotted against *rms slope* in *Figure 3.3.2*. As we would expect, the accuracy of the approximation falls with increasing slope—falling below the 10 dB mark for surfaces with an rms slope greater than 0.25.

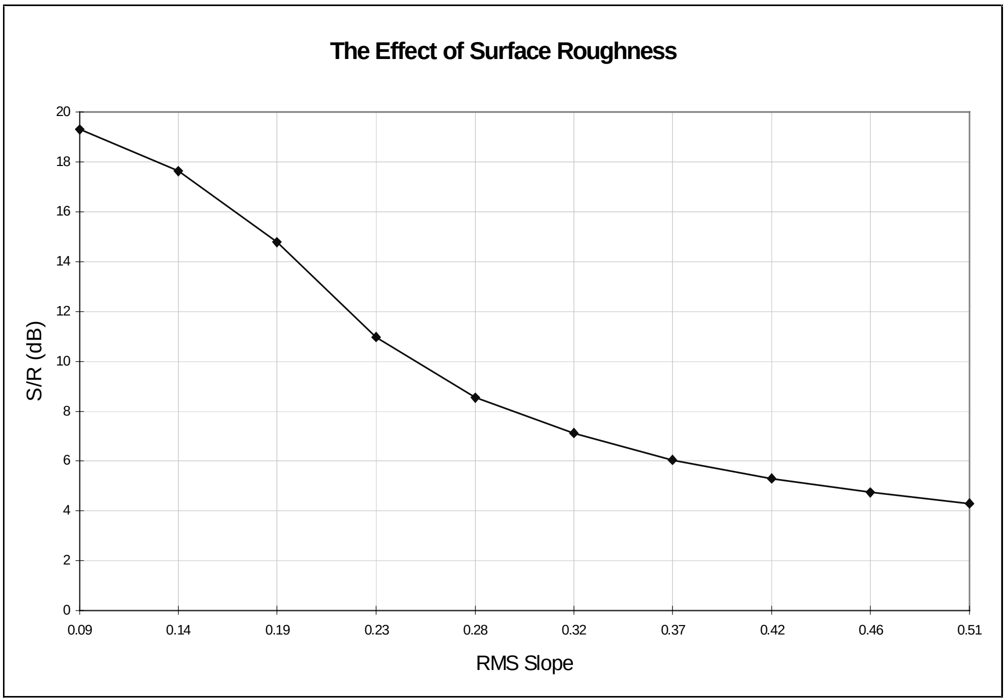

Figure 3.3.2 The effect of surface slope on the accuracy of an optimal filter.

Surface Directionality

If we consider Ogilvy's directional surface and alter the tilt of the illuminant vector we see that the S/R of the prediction is highly dependent on the tilt of the illuminant relative to the directionality (*Figure 3.3.3*). The prediction is least accurate when the illuminant vector is parallel with the grain of the surface, falling to a minimum of 5 dB. We also note that the increasing rms roughness in the principal direction causes a significant decrease in the prediction accuracy.

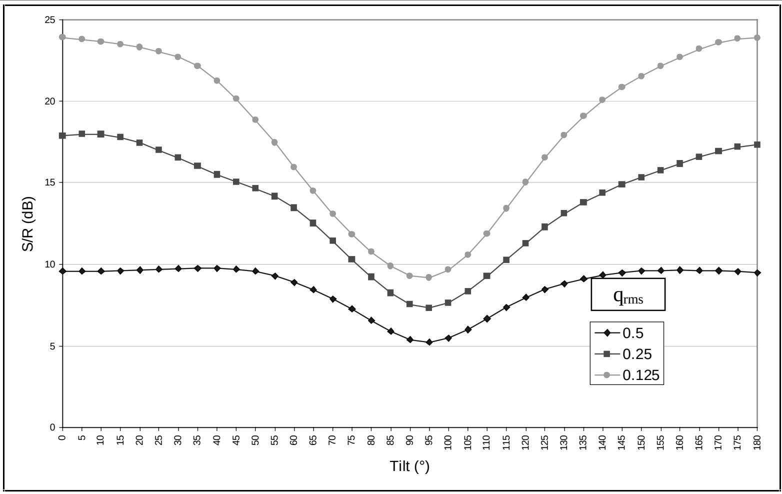

Figure 3.3.3 The accuracy of image prediction for a directional surface.

If we again set an arbitrary signal to noise limit of 10 dB we may form the following nominal bounds for the application of a linear approximation: surfaces should have an rms gradient not exceeding 0.25, whereas strongly directional surfaces should not be illuminated from tilt angles at less than 15 ° to the perpendicular of their primary axis.

3.3.4 Summary

In this section we found that the accuracy of the approximation falls as *rms slope* increases. Setting a nominal lower limit of 10dB for the S/R ratio we showed that isotropic surfaces could be modelled as long as the *rms slope* did not exceed 0.25. The accuracy with which directional surfaces could be modelled is highly dependent on the angle between the grain of the material and the illuminant. Again taking the 10dB limit, we found that surfaces with *rms slopes* of less than 0.25 could be accurately modelled except in the region ±15° of the material grain direction.

3.4 Kube's Model

Having considered the accuracy and scope of an optimal linear filter, in this section we shall consider an analytical linear model reported in the literature [Kube88]. Since the characteristics of the function, which the optimal model approximates vary with surface characteristics, a single analytical filter cannot be optimal for all possible surfaces. The *form* of the model may, however, be common to the optimal filters of a range of surfaces. The next question we ask is under what circumstances does the form of the theoretical model coincide with the behaviour of the least squares filter? Kube's model predicts that the intensity will be a linear combination of the surface derivatives scaled by a trigonometric function of tilt as well as a mean component. If the coefficients of the optimal linear model, recalculated for each tilt condition, vary in the same manner as the trigonometric function, we may conclude that the form of Kube's model is optimal in the least squares sense.

3.4.1 An Analytical Expression

We restate Kube's model in our own notation, dropping the assumption that the surface is fractal.

Consider the bandlimited surface with height map defined by the scalar function $s(x,y)$.

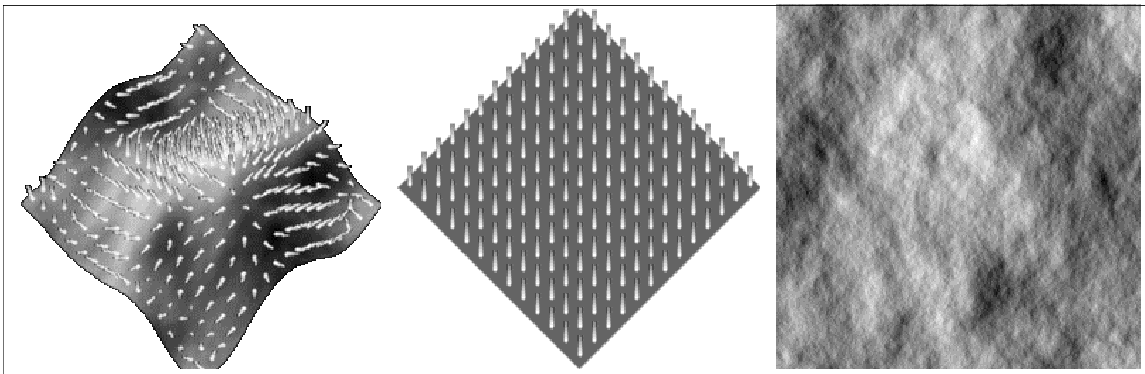

(a.) **S(**x,y**)** *(b.)* **L(**x,y**)** *(c.) i(x,y)*

Figure 3.4.1 The components of image formation.

(a) The Lambertian Image

let **S(x,y)** (*Figure 3.4.1*a) be the derivative field of the bandlimited scalar field *s(x,y)* such that

$$\mathbf{S(x,y)} = \quad \text{grad } s(x,y) = \begin{bmatrix} \frac{\delta z}{\delta x} & \frac{\delta z}{\delta y} & \frac{\delta z}{\delta z} \end{bmatrix}$$

Also define the illumination vector field **L(x,y)**, assuming illumination is produced by a point source located an infinite distance from the scene, the magnitude and direction of the vectors will be uniform throughout the field (*Figure 3.4.1*b).

$$\mathbf{L(x,y)} = \begin{bmatrix} k_1 & k_2 & k_3 \end{bmatrix}$$

where

$k_1 = Cos\tau Sin\sigma$ $\quad$ $k_2 = Sin\tau Sin\sigma$ $\quad$ $k_3 = Cos\sigma$

Assume that a surface has a reflectance function which is:

(i) homogeneous over the surface, and

(ii) Lambertian.

Furthermore, assume that the surface is constrained such that there is no significant:

(i) cast or self shadows, or

(ii) degree of interreflection.

Adopting these assumptions we may state that the image field *i(x,y)* (*Figure 3.4.1*c)is the normalised scalar product of the surface derivative field and the illumination vector field.

$$i(x,y) = \frac{\mathbf{L}(x,y)^T \mathbf{S}(x,y)}{\left|\mathbf{L}(x,y)^T \mathbf{S}(x,y)\right|}$$

This results in a non-linear operation at each facet,

$$i = \frac{-k_1 p - k_2 q + k_3}{\sqrt{p^2 + q^2 + 1}} \qquad (3.4.1a)$$

where i, *p* and *q* are functions of *x* and *y*.

(b) Kube's Linearisation

Kube uses the Taylor Expansion to form a linear approximation to equation 3.4.1, shown below, truncated beyond the quadratic term.

$$\text{i}= k_1 p + k_2 q + k_3 - \frac{k_3}{2}\left(p^2 + q^2\right)$$

Using the first three terms he forms a linear estimate

$$\hat{i} = k_1 p + k_2 q + k_3 \qquad (3.4.1b)$$

$$\text{or} \qquad \hat{i} = \begin{bmatrix} k_1 & k_2 & k_3 \end{bmatrix} \begin{bmatrix} p \\ q \\ 1 \end{bmatrix}$$

This approximation is reasonable where:

$p >> p^2$ and $q >> q^2$

i.e. where

$p << 1$ and $q << 1$

(c) Frequency Domain Dual

Since equation 3.4.1b is linear, we can easily form a frequency domain dual. For simplicity we will work in terms of the magnitude (denoted by the *m* subscript rather than the power spectrum.

$$I_m(u,v)=P_m(u,v).k_1+Q_m(u,v)k_2+k_3\delta(0,0)$$

where $P_m(u,v)$ and $Q_m(u,v)$ are the spectra of the directional derivative fields.

However, the directional derivatives are related to $S(u,v)$

$$P_m(u,v) = iuS_m(u,v)$$
$$Q_m(u,v) = ivS_m(u,v)$$

where *i* represents a 90° phase shift.

If we restate these in polar co-ordinates:

$$P_m(\omega,\theta) = i\omega\cos\theta.S_m(\omega,\theta)$$
$$Q_m(\omega,\theta) = i\omega\sin\theta.S_m(\omega,\theta)$$

Ignoring the mean and using polar co-ordinates:

$$I_m(\omega,\theta) = k_1.i.\omega.\cos\theta.S_m(\omega,\theta) + k_2\ i.\omega.\sin\theta.S_m(\omega,\theta)$$

Resolving the k terms into their trigonometric components and gathering like terms:

$$I_m(\omega,\theta) = i.\omega.\sin\sigma.S_m(\omega,\theta).[\cos\theta.\cos\tau + \sin\theta\sin\tau] \qquad (3.4.1c)$$

Which can be simplified to

$$I_m(\omega,\theta) = i.\omega.\sin\sigma.\cos(\theta-\tau).S_m(\omega,\theta)$$

From the point of view of this thesis, the most important effect of this expression is that the image spectrum is a function of both the surface and the illuminant tilt.

3.4.2 A Comparison of Kube's Model with the Optimum Linear Model

Kube's model gives an analytical expression as to the relationship between surface and image, however, we can also empirically define a least squares mapping from *p* and *q* to intensity as we did in section 3.3.1.

We can therefore define a least squares model, with parameters *a* and *b*, to map the surface derivative fields to the image intensity fields for a given tilt. If we then vary tilt we can compare the behaviour of the parameters (*a* and *b*)with that predicted by Kube, i.e. $k_1=\cos\tau.\sin.\sigma$ and $k_2=\sin\tau.\sin.\sigma$ (σ is held constant at 60° in this work).

Consequently if Kube's model is near optimal in the least squares sense we would expect that the *a* and *b* parameters of the linear model to be equivalent to the parameters k_1 and k_2 and would follow $\cos\tau$ and $\sin\tau$ curves respectively as tilt is varied and the slant angle is held constant. We therefore plot the estimates of *a* and *b* for an isotropic surface against these expected functions in *Figure 3.4.2*.

If we test this prediction with an isotropic surface we see a linear relationship between the parameters estimated by the least squares criterion and those predicted by Kube's equation *3.4.1c*. However, it is immediately obvious that the slope of the line varies with the rms slope of the surface; this was not predicted by Kube's model. If we tabulate these results we can see that both the estimated parameters are highly correlated with Kube's predictions (see *Table 3.4.1*). It is also clear that the scaling factor falls as *rms slope* increases. Finally we note the estimated a and b parameters at a given tilt are scaled by almost equal amounts. We measure the degree of correlation between the parameters *a* and *b* with Kube's predictions.

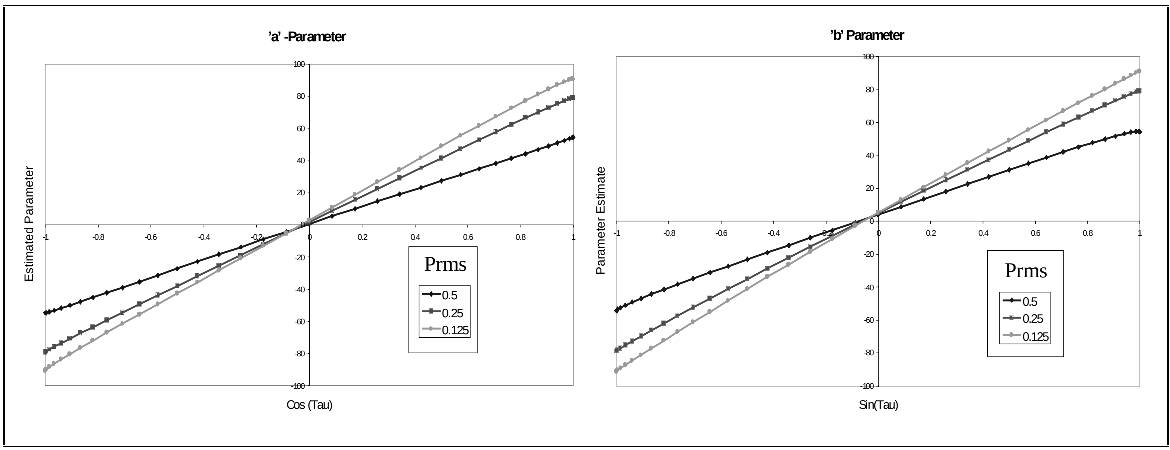

Figure 3.4.2 The variation of the least squares parameters plotted against those predicted by Kube (isotropic fractal surface).

	Scaling		**Correlation**	
RMS Slope	**a**	**b**	**a**	**b**
0.125	90.83	90.89	0.99988	0.99963
0.250	79.21	78.89	0.99995	0.99956
0.500	54.60	54.37	0.99995	0.99945

Table 3.4.1 The relationship of LS parameters to Kube's predictions.

If we now consider Ogilvy's directional surface we note the relationship is less linear (*Figure 3.4.3*), though the correlation coefficients are still high (*Table 3.4.2*). The trend of decreasing line slope with rising rms surface slope is still evident, though there is a growing disparity between the parameter scaling values as slope increases. This would seem to be in agreement with the inverse slope relationship found in both the isotropic and directional cases.

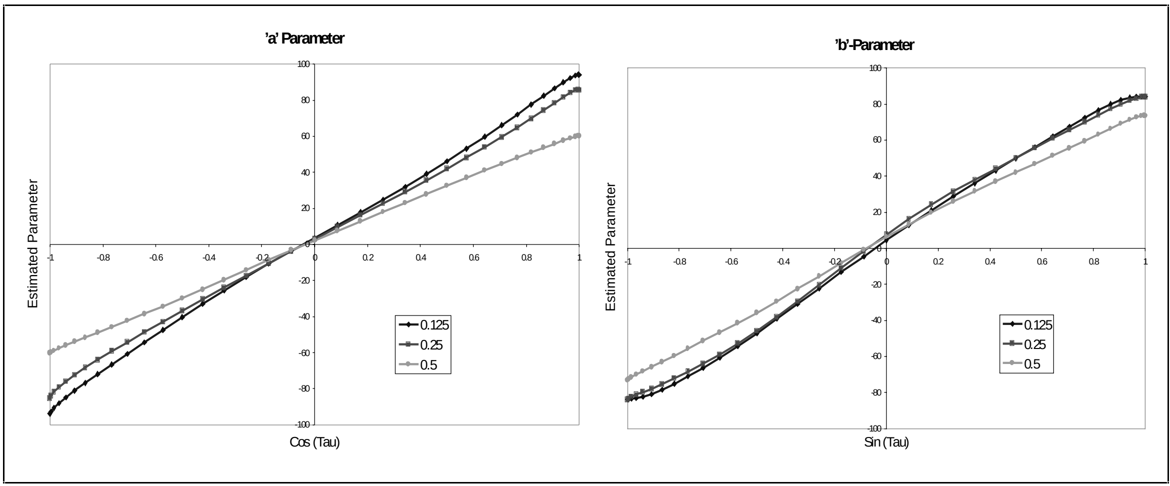

Figure 3.4.3 The variation of LS parameters for directional Ogilvy surface to Kube's predictions.

Slope		**Scale**		**Correlation**	
$\mathbf{p_{rms}}$	$\mathbf{q_{rms}}$	**a**	**b**	**a**	**b**
0.125	0.34	91.72	89.46	0.999537	0.9985
0.250	0.067	82.90	87.77	0.999554	0.9983
0.500	0.134	60.75	74.64	0.999833	0.999217

Table 3.4.2 The relationship of LS parameters for directional surface to Kube's predictions.

3.4.3 Summary

We conclude that Kube's model is near optimal in the least squares sense for isotropic surfaces, albeit with the proviso of a scaling factor dependent on the *rms slope* of the surface. Furthermore, while some distortion is apparent in the parameter estimates, Kube's model agrees to a satisfactory degree for Ogilvy's directional surfaces as long as the slope is moderate. The scaling effect is also apparent for the directional surface.

3.5 Real Textures

3.5.1 Test textures and their Spectra

In this section we shall use two exemplar textures and will consider whether the predictions of Kube's model and our simulations are borne out by real data. In this chapter we will concentrate on two textures: the Rock texture, *Figure 3.5.1*, and the Striate texture, *Figure 3.5.2*. The Rock texture is isotropic and of the textures used in this thesis, is the most similar in appearance to the fractal simulations. The Striate texture is highly anisotropic, and is the closest experimental texture to the synthetic Ogilvy surface.

Figure 3.5.1 Rock Texture

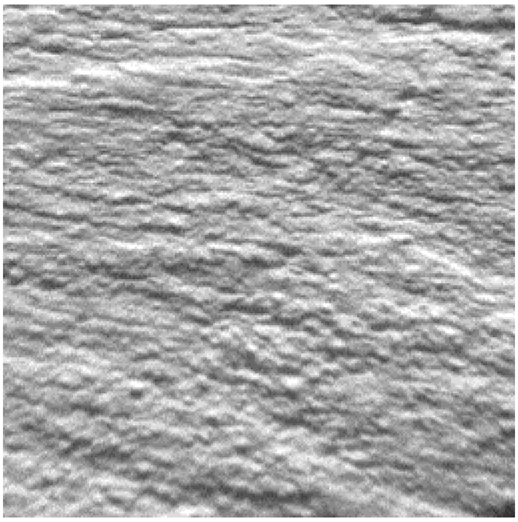

Figure 3.5.2 Striate Texture

If we consider the log magnitude spectra of the above images (*Figure 3.5.3*), we note a valley running at right angles to the direction of illumination in both spectra. This is consistent with Kube's prediction of attenuation of frequency components at right angles to the illuminant direction. In addition, the anistropic texture has an obvious ridge running in the perpendicular direction which is not apparent in the isotropic case.

Figure 3.5.3 Magnitude Spectra of Rock (left) and Striate (right) Textures τ=0°

If the light source revolves through 90° the power spectrum of the isotropic surface is effectively rotated by the same amount, however, the change in tilt causes the ridge of the striate texture to be lost. While the effect on the isotropic texture is equivalent to one of rotation, the image spectrum of the directional surface is of a fundamentally different image texture.

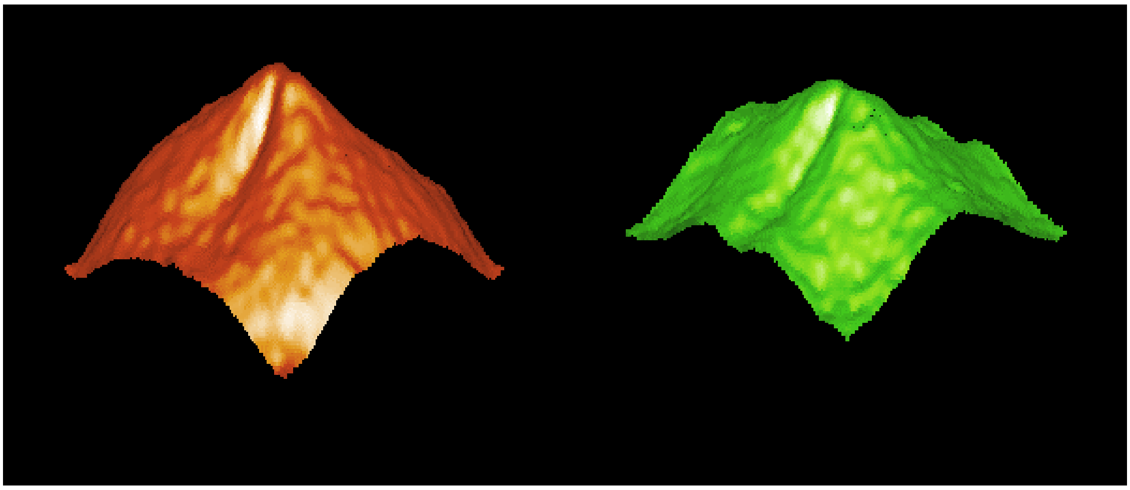

Figure 3.5.4 Magnitude Spectra of Rock (left) and Striate (right) Textures τ=90°

3.5.2 The Polar Power Distribution

In order to verify Kube's model, however, it is necessary to consider the effect more analytically. We focus our attention on the variation of spectrum power with polar angle. This is carried out by integrating the power along radii of the polar spectrum and plotting the quantity as a function of polar angle. Due to the finite quantisation of frequency at the low end of the spectrum, and the effects of noise at the high end, this integration is performed within certain frequency bounds. The bounds used for our first experiment are $\omega = 0.156f_s$ to $0.469f_s$ shown in *Figure 3.5.5*.

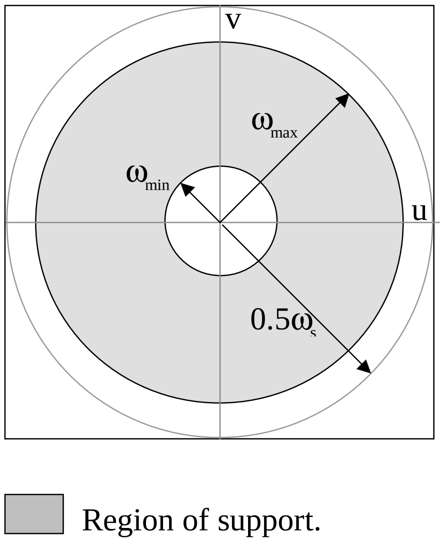

Figure 3.5.5 Region of support for polarogram.

Plotting this quantity for the rock texture *Figure 3.5.6* with the least squares fit of both the Tau 0 and 90° images, we find that it is in agreement with the $|\text{Cos}(\theta-\tau)|$ (or $\text{Cos}^2(\theta-\tau)$ for the power spectrum) relationship predicted from Kube's work with the addition of a small bias term. This result is also in agreement with that reported by Chantler [Chantler94] for a similar isotropic rock surface.

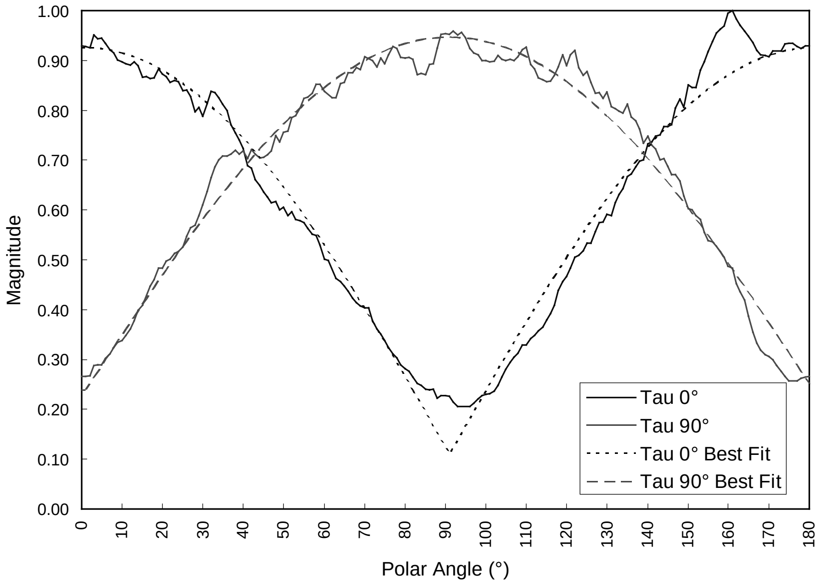

Figure 3.5.6 The variation of image magnitude with polar angle.

In the case of an isotropic surface, the dominant source of image directionality is due to directional illumination. However, if the surface is anisotropic, the surface directionality will also contribute to the polar distribution of power, *Figure 3.5.7*.

Figure 3.5.7   Polar distribution of image power for Striate texture.

3.5.3 Summary

In this section we took two real textures and showed that their image spectra are directionally attenuated as predicted by Kube's model. Using the isotropic test texture it was possible to verify the $|\cos(\theta-\tau)|$ relationship predicted by Kube. For the isotropic surface, varying illuminant tilt is equivalent to rotating the texture, however, for the directional surface the attenuation causes a fundamental change in the visual texture which cannot be modelled as rotation.

3.6 Conclusions

This chapter has considered the underlying mechanisms of reflection and discussed four analytical descriptions of the phenomenon. Three of the models were implemented and compared with the function measured for the surface type used in the experimental work. Both Oren's and Lambert's models were found to be effective methods of modelling, and, whereas Oren's model was the more accurate over a range of incident angles, Lambert's Law was found to be more accurate when the slope distribution of a rough surface was taken into account. We then considered the circumstances under which a linear filter could maintain adequate accuracy. Taking a nominal accuracy threshold of 10dB we concluded the linear form could accurately predict images of surfaces with an rms slope of less than 0.25. The highly directional test surface had the additional restriction that the illuminant should not make an angle of less than 15° from the grain direction of the material.

The third part of this chapter considered an analytical frequency based description of the link between surface and image. A model reported in the literature, [Kube88] was investigated to see whether it is optimal in the least squares sense. We concluded that the form was optimal for most surfaces, provided that a scaling factor dependent on the surface roughness was used.

The observation of the real isotropic surface *Rock* confirmed Kube's predictions for the angular distribution of power. The behaviour of the directional surface *Striate* illustrated that the tilt effect can either attenuate or accentuate useful directional information. This implies that tilt variation will have serious consequences for the discrimination of rough surfaces.