The effect of variation in illuminant direction on texture classification

Michael J Chantler

Thesis submitted
for the
Degree of Doctor of Philosophy

Department of Computing and Electrical Engineering
Heriot-Watt University

August 1994

Table of contents

List of figures

List of tables

Acknowledgements

I would especially like to thank my supervisor, Professor George Russell, for his support, ideas, encouragement, and stimulating discussion.

Special thanks are also due to Laurie Linnett for his ideas, time, enthusiasm, and for the software he selflessly provided.

I would like to thank all the other members of the department for their help, particularly Stuart Clarke, Andy Wallace, Keith Brown, Alastair Houstin, Matt Dunnigan, Mary Kiernan, Ceri Reid, Andrew Quinn, Phil Knightbridge, David Lindsay, and many others who have helped me in my work.

Lastly, I would like to say that I am especially indebted to my wife Linda for all her support and backing.

Abstract

Texture analysis has been an extremely active and fruitful area of research over the past twenty years. Many advances have been made, but the effect of variation in lighting conditions on automated texture classification and segmentation has received little attention. This thesis shows that the direction of the illuminant is an important factor that should be taken into account when analysing images of three-dimensional texture.

A frequency domain model is presented which predicts that both the directional characteristics and the variance of images of three-dimensional texture can be affected by changes in illuminant vector. Results of simulations and laboratory experiments support these predictions.

The responses of three sets of texture measures are analysed using a test set of isotropic and directional textures. The results show that the feature measures' outputs are affected by changes in illuminant direction. These changes are also shown to significantly increase the error rates of statistical classifiers implemented using the three feature sets. Normalisation of images is shown to reduce the error rates in some cases.

The frequency domain model of image texture is further developed using empirical data and the resulting model used to design a set of tilt-compensation filters. These filters are used to pre-process images to reduce the effects of changes in the angle of tilt of the illuminant. Application of the filters to the test image set reduced the classification errors associated with directional textures.

Principal symbols and abbreviations

Symbol	Meaning	Section first introduced
σ_S	angle of slant of the surface	2.1.3
τ_S	angle of tilt of the surface	2.1.3
$\mathcal{F}[g(x,y)]$	the Fourier transform of the function $g(x,y)$	2.2.1
σ	angle of slant of the illuminant vector **L**	2.2.1
τ	angle of tilt of the illuminant vector **L**	2.2.1
L	unit vector pointing at the light source	2.2.1
n	unit vector normal to the surface	2.2.1
$I(x,y)$	normalised intensity image	2.2.1
$F_I(\omega,\theta)$	Fourier transform of the normalised intensity image	2.2.1
$V_H(x,y)$	surface height map	2.2.1
$F_H(\omega,\theta)$	Fourier transform of the surface height map	2.2.1
$F_s(\omega,\theta)$	surface response component of $F_H(\omega,\theta)$	2.2.1
$F_\tau(\omega,\theta)$	tilt response component of $F_H(\omega,\theta)$	2.2.1
$F_\sigma(\omega,\theta)$	surface response component of $F_H(\omega,\theta)$	2.2.1
ω	angular frequency (polar co-ordinates)	2.2.1
θ	angle of direction w.r.t. the x-axis (polar co-ordinates)	2.2.1
β_H	surface power roll-off factor	2.2.1
β_I	image power roll-off factor	2.2.1
S	height scaling factor	3.1.1
s_H^2	surface height variance	3.1.3
s_H	surface height deviation	3.1.3
$\hat{\alpha}$	average estimated slope angle	3.1.3
b_τ	y-intercept parameter of the raised cosine model of the tilt response component $F_\tau(\omega,\theta)$	3.3.4(c)
m_τ	slope parameter of the raised cosine model of the tilt response component $F_\tau(\omega,\theta)$	3.3.4(c)
ω_1	angular frequency in the x direction	5.1
ω_2	angular frequency in the y direction	5.1
$H(\omega_1,\omega_2)$	transfer function	5.1
$P(i,j)$	un-normalised co-occurrence matrix	5.2
$p(i,j)$	normalised co-occurrence matrix	5.2
d	co-occurrence matrix displacement vector	5.2
D^2	Mahalanobis distance	5.4
$\overline{D_\tau^2}$	mean tilt sensitivity	5.4
ω_s	sampling frequency	6.4.1

Abbreviation	**Meaning**	**Section first introduced**
ARMA	autoregressive moving average (model)	4.3.3
ASM	angular second moment (co-occurrence feature)	5.2
CON	contrast (co-occurrence feature)	5.2
cooc1	feature set/classifier based on co-occurrence features	6.2.2
COR	correlation (co-occurrence feature)	5.2
ENT	entropy (co-occurrence feature)	5.2
FBM	fractional Brownian motion	2.1.3
frac1	feature set/classifier based on Linnet's features	6.2.2
GLCM	grey-level co-occurrence matrix	4.4.1
laws1	feature set/classifier based on Laws' features	6.2.2
MRF	Markov random field	2.1.1
PCA	principal components analysis	4.3.1
PSD	power spectral density	2.2.1
rms	root mean square	2.1.4.
ROV	remotely operated vehicle	1.1
TEC	total error of classification	6.1.1

Chapter 1

Introduction

1.1. Motivation

The original motivation for the work described in this thesis stemmed from a desire to segment underwater video images taken by remotely operated vehicles (ROVs). It was thought that the ability to interpret these images would, in conjunction with the processing of data from range sensors, enable simple vision tasks to be undertaken; such as pose determination of the cylindrical components of underwater structures. Unfortunately the sub-sea environment is an extremely hostile one in which to attempt such tasks. Back scatter from plankton and other suspended matter, marine growth, corrosion, and the fact that the majority of underwater structures are painted dark grey, mean that images are often noisy and of poor contrast. However, the different textures in the image, caused by back scatter and marine growth etc. may be exploited. Application of a texture segmentation and classification scheme [Linnett91a] to images of an underwater installation produced good results [Chantler91]. This raised the question of what would happen to the appearance of the textures when an ROV moved around a structure — as the position and orientation, of both the viewer and the vehicle mounted illumination, would change relative to the physical texture.

As the research proceeded the scope of the work was reduced to that of investigating the effects of changes in illuminant direction. Such variations may be encountered in a variety of situations. Close proximity point lighting, often used for inspection purposes, provides illumination at varying angles over the scene. Remote sensing devices sensitive to the visible spectrum experience variations in illuminant vector according to the time of day. Many other remote sensing systems that provide their own illumination, e.g. active sonar and radar, are non-stationary and hence the illuminant vector is dependent upon the approach and orientation of the survey platform. Thus there are a wide range of applications in which texture classification may have to be performed under varying illumination conditions. The work described here was therefore divorced from the original underwater application.

Thus the aims of this thesis are :

(i) to provide an understanding of the effects that variation in illuminant direction has on images of physical texture,

(ii) to investigate the impact of these effects on texture classification and segmentation, and

(iii) to propose methods of reducing classification errors caused by variation in illuminant direction.

1.2. Scope of the research

This section outlines the scope of the work described in this thesis. For reasons of brevity not all of the restrictions are described here — further details are given in the appropriate chapters.

Texture classification, as referred to in this thesis, normally involves three processes, as illustrated in figure 1.1.

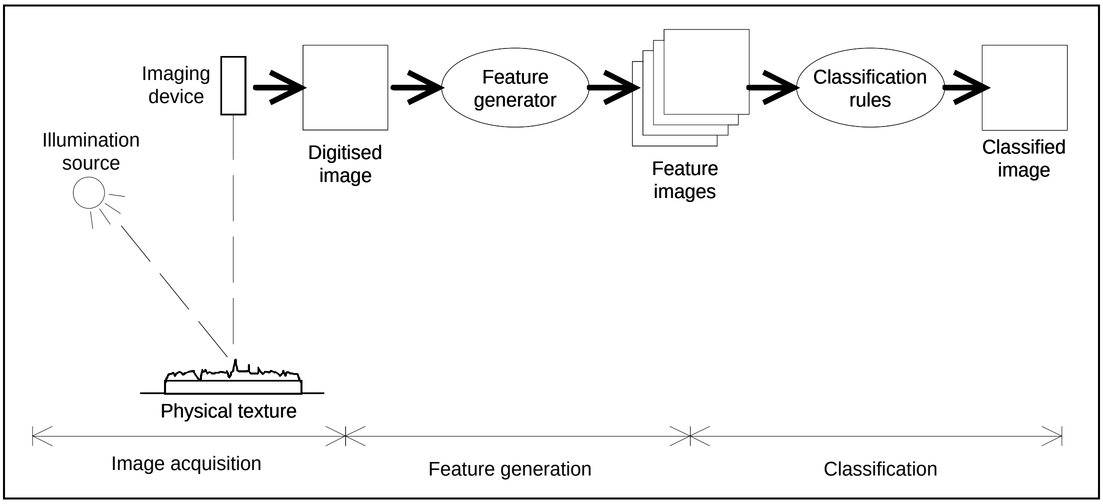

Figure 1.1 - Processes in texture classification

First, the subject texture must be illuminated and its image acquired (in the case of images taken in the visible spectrum this is normally performed using a stills camera and scanner, or alternatively a video camera and frame store). Second, feature operators are applied to the digitised image to produce a set of feature images that provide a numerical description of the characteristics of the texture(s). Third, a set of rules is applied to classify the image into texture classes.

This thesis concentrates on the effects of variations in the image acquisition process. More specifically it is concerned with the effects that changes of illuminant direction have on feature generation and classification.

Among other restrictions, the illumination is assumed to be unidirectional and the physical texture assumed to consist only of surface relief; that is it is assumed to contain only *topological texture*[1]. The experiments are confined to the variation of illuminant direction; the viewer's position being fixed vertically above the physical texture which lies upon a horizontal plane (as depicted in figure 1.1). The investigation into the effects on classification is restricted to the case where the illumination is varied between training and classification sessions.

1.3. Thesis organisation

This thesis essentially consists of two parts. The first part, comprising chapters 2 and 3, is concerned with the image acquisition process; that is it investigates the effects of variation in illuminant direction on image texture. The second part, comprising chapters 4, 5, and 6, examines the impact that these effects have on feature generation and classification, and proposes methods for reducing the resulting errors.

Hence chapter 2 provides a short review of research into the effects of illuminant variation on image texture. One model due to Kube and Pentland [Kube88] is identified and presented. Its implications for texture classification are assessed. Chapter 3

---

[1]The term *topological texture* is used solely to refer to the three-dimensional variation, or relief of a physical surface. In contrast the term *albedo texture* is used to refer only to *surface markings. Image texture* consists of intensity variations in the image plane and can be due to either topological or albedo texture or a combination of the two. However, as stated above, only the former is of direct concern here.

investigates the validity of this theoretical model using simulations and laboratory experiments.

Having investigated the effect of changes in the illumination direction on image acquisition, the second part of the thesis addresses its effect on feature generation and classification. In chapter 4 feature sets employed in texture classification are surveyed and three are selected for further investigation. Chapter 5 reports the results of this further investigation which uses images of isotropic textures captured under controlled illumination conditions. In chapter 6 these same images, augmented with images of a directional texture, are made up into montages and used to test the effects of variation in illuminant vector on three classifiers. The second half of this chapter proposes several methods for the reduction of classification errors induced by changes in illuminant tilt[2], and one proposal is implemented and tested on montages of the test textures.

Finally, in chapter 7, the work is summarised and final conclusions presented.

1.4. Original work

It is believed that this thesis contains two topics which represent original work. First, the effect of variation in illuminant direction on supervised texture classification has been explicitly investigated; and second, a compensation scheme has been developed which is shown to be capable of reducing classification errors induced by changes in illuminant tilt.

1. In chapter 5 variation of the illuminant vector is shown to affect the outputs of three sets of texture features when they were applied to images of isotropic textures, and an existing metric — the Mahalanobis distance — is adapted for use as a new measure of sensitivity to tilt variation. In chapter 6 classifiers based upon these three feature sets are shown to be adversely affected by changes in illuminant direction between training and classification sessions. While this

[2]The tilt angle of the illuminant as referred to here, is the angle that the projection of the illuminant vector onto the texture reference plane makes with an axis in that plane. Its companion, slant angle, is the angle that the illuminant vector makes with a normal to the reference plane (see figure 2.1).

behaviour seems obvious, it is believed that it has not been investigated and reported before from the perspective of automated texture classification.

2. The second half of chapter 6 is devoted to the development of a scheme that is designed to compensate for variation in illuminant tilt. This scheme comprises a set of filters derived from a frequency domain model of image texture. The model, originally due to Kube and Pentland [Kube88], was further developed in the light of empirical evidence presented in chapters 3 and 6. Application of the scheme to three classifiers reduced tilt induced errors when tested on montages of isotropic and directional textures. It is believed that this scheme represents a novel approach to illuminant tilt compensation.

Chapter 2

Image models of topological texture

This chapter surveys possible sources of mathematical models or empirical studies that would enable predictions to be made about the effect of illuminant vector variation on images of topological texture. It reviews models of image texture in general, and places particular emphasis on models that define image characteristics in terms of topological texture and illumination parameters. The latter type will be referred to as *image models of topological texture.* This terminology is necessary in order to differentiate such models from the purely two-dimensional *image texture models,* used extensively by the texture classification community, and the *"three-dimensional texture models"* — models of albedo texture on three-dimensional surfaces used by computer graphics and shape from texture researchers [Cohen91c] [Patel91].

Thus the term *image model of topological texture* is exclusively used to refer to a model that, given certain characteristics of the physical surface together with a description of the illumination and the viewer's position, can be used to predict characteristics of texture in the image.

Four areas of research would seem to be likely candidates for the development of such models :

(i) texture synthesis (mainly used in computer graphics to add realism to images),

(ii) texture segmentation and classification,

(iii) shape from texture, and

(iv) scattering theory.

Each of these areas will now be reviewed in turn with the objective of identifying sources of suitable theory or empirical studies, that will enable predictions to be made about the behaviour of texture under varying illumination conditions. These reviews are brief — as there is surprisingly little in the way of published literature on the effects of lighting on image texture in the first three areas, while the last is not directly applicable. These review sections are followed by an examination of one model in detail, and the chapter concludes by considering the implications that this model has for texture analysis.

2.1. Review

2.1.1. Texture synthesis

One of the most frequent criticisms of early computer graphics was the lack of realism due to the apparent smoothness of the three dimensional surfaces portrayed. It is not surprising therefore that one of the main uses of texture synthesis has been to improve the realism of such graphics. Heckbert [Heckbert86] surveyed texture mapping techniques which are concerned with mapping *two-dimensional* arrays or functions of texture onto screen space according to the three-dimensional surfaces contained in object space. Texture mapping is most commonly used to modulate surface colour [Blinn90] and for "bump mapping" i.e. surface normal perturbation [Blinn78] [Haruyama84] [Baston75]. The former treats texture purely as a set of surface markings, while the latter provides a simplified way of imitating the effects of topological texture (occlusion and shadowing are ignored). Blinn [Blinn90] states that surface marking based schemes produce images that look like smooth surfaces with photographs of wrinkles glued on — as the *light source directions* are rarely the same in the original texture map and graphics model. He notes that the effect of wrinkles on intensity is primarily due to variation of the surface normal, and therefore goes on to develop a texture scheme based upon *small perturbations* of surface normals. His results are extremely realistic, and justify his assumption that the major effects of topological texture (consisting of small perturbations) can be modelled solely as variations of the surface normal. This assumption is also made in a later section of this chapter, which presents an image model of topological texture due to Kube and Pentland [Kube88].

An alternative to texture mapping was first developed by Gagalowicz and Ma [Gagalowicz86]. Their model-based approach essentially parameterises a planar texture model (based on second order statistics) with three-dimensional spatial parameters.

Synthesis is thus performed directly on the surface and avoids the mapping process above. Cohen and Patel developed a similar model-based approach but they used the more parsimonious Markov random field (MRF) model [Cohen91c] [Patel91] [Patel93]. Their "three-dimensional texture model" is a two-dimensional MRF model, with additional surface shape parameters, that enables the foreshortening effects due to surface orientation and perspective projection to be taken into account. Neither of these approaches are however of direct interest to this survey, as they do not consider topological texture or illuminant effects.

A very popular area of computer graphics that does use models of topological texture, and does take illumination into account, is that of fractals [Mandelbrot85]. Spectacular "natural" images have been generated by Voss [Voss88], Saupe [Saupe88], Bouville [Bouville85] and others, using random fractals. These researchers are primarily concerned with the appearance of the final image, and while they do use stochastic topological texture models, and do explicitly take into account illumination, they have not in general developed corresponding image models. Pentland [Pentland84] [Pentland86] in his shape from shading work did however investigate such a model, and this is discussed in the *Shape from texture* section of this review.

Boulanger, Gagalowicz, and Rioux [Boulanger89] also used topological texture models – but for data compression purposes. They recreated the appearance of surface texture on museum artefacts using an autoregressive model and Lambertian shading model, but, as for the majority of the fractal work, they were primarily concerned with the appearance of the resulting image and its image model was therefore not investigated.

To summarise : texture synthesis researchers have explicitly considered and used models of topological texture (e.g. fractals) and have taken into account lighting conditions. However, as their primary concern is the appearance of the final image, they have no requirement or motivation to develop mathematical models of the resulting image texture. Cohen et al, and Gagalowicz et al, developed texture models that incorporate surface orientation and camera projection parameters. However, they did not take lighting effects into account and their texture models are two-dimensional.

2.1.2. Texture analysis - segmentation and classification

Random field models have been used for the synthesis of perfect test textures with known and consistent characteristics for the testing of texture segmentation and classification algorithms. The models themselves have also been used as the basis of texture segmentation and classification methods. If a model is capable of representing and synthesising a range of textures, then estimates of its parameters may provide a useful feature set. For such a model-based approach to be successful there must exist a reasonably efficient and appropriate parameter estimation scheme and the model itself should be parsimonious, i.e. use the minimum number of parameters. Popular random field models used for texture analysis and testing include fractals [Pentland84] [Peleg84] [Medioni84], autoregressive models [Kashyap80] [Khontanzad87], fractional differencing models [Kashyap84] [Choe91a], and Markov random fields [Chellappa85a] [Cohen91b]. To the best of the author's knowledge (see chapter 4 for a detailed review) all of these models are used purely as image texture models; that is they are used to represent and synthesise two-dimensional intensity textures directly in the image plane. Only very rarely is consideration given to topological texture and lighting. Indeed, even on the wider subject of machine vision, few papers or books give details of lighting schemes used for image acquisition, and fewer still give any background theory for such schemes [Davies90].

Davis [Davis81a] describes two approaches to modelling image texture. A *"physically based"* model takes into account surface relief, albedo, illumination, and the position and frequency response of the viewer. *"Image-based"*, models on the other hand, model textures directly in the image plane without regard to their physical origin. He states that physically based models are ordinarily very difficult to construct, and this in part explains their scarcity. Davis however, does not consider the effects of illumination in his experiments. Nor do many of the other papers on texture segmentation and classification (see chapter 4). This is particularly surprising for papers on *"rotation invariant"* schemes, where one might reasonably expect researchers to rotate the physical textures on their own without rotating the associated lighting.

Pentland is one of the few researchers to develop a segmentation scheme who does consider topological texture and lighting [Pentland84]. He states that *"the lack of a 3-D model for such naturally occurring surfaces has generally restricted image-understanding efforts to a world populated exclusively by smooth objects, a sort of 'Play-Doh' world".* He uses fractal dimension as a feature measure, and shows that it is theoretically independent of illuminant direction. That is he proves that the fractal dimension of texture in the image plane, is the same as the fractal dimension of the components of the normals of the physical surface being imaged (assuming a Lambertian reflectance function, constant illumination and constant albedo).

Pentland also used fractal models in his shape from texture algorithms, and these are described in the next section.

2.1.3. Texture analysis - shape from texture

This section gives a brief overview of the two main *shape from texture* techniques — texture gradient based approaches and isotropy based approaches. This is followed by a more detailed review of the associated literature that has considered topological texture and illumination issues. *Shape from shading* techniques - e.g. [Ikeuchi81], [Horn89]- have not been reviewed here, as they normally assume that the surfaces under consideration are smooth [Pentland86], and they do not employ models of topological texture.

a) Texture gradient and isotropy approaches

There are two ways that surface shape affects images of texture. Firstly, perspective projection effects a uniform compression which is dependent on the distance of the surface from the viewer, the greater the distance the greater the compression. Secondly, projection of surfaces that are not perpendicular to the viewing direction will result in a foreshortening effect. The degree of foreshortening is proportional to the cosine of the surface inclination angle (surface slant angle σ_S), and the direction of maximum foreshortening is the direction of the steepest descent (surface tilt angle τ_S).

Correspondingly, two approaches have been employed to estimate the tilt and slant of surfaces. The first exploits the concept of "texture gradients" and is due to Gibson

[Gibson50]. It assumes that physical texture is homogeneous and exploits the gradient of texture densities caused by perspective projection. The second uses a statistical approach first proposed by Witkin [Witkin81]. The distribution of orientations in a texture image is biased towards a direction perpendicular to the tilt angle and the degree of biasing is a function of the slant. Both surface slant and tilt can therefore be estimated from the distribution of orientations (assuming that the original texture is isotropic). These two approaches have been extensively researched. Bajcsy [Bajcsy76] uses a texture gradient based on "preferred" frequencies derived from Fourier transforms of 128x128 windows. Blostein [Blostein89] expicitly identifies texture elements (*textels*) in textures. She defines a *texel* as "*the repetitive unit of which the texture is composed*", and uses the *texel* area gradient to extract depth information. Rosenfeld [Rosenfeld75] suggested the use of an edge operator as a simple method of measuring texture gradient. Researchers who have built upon Witkin's ideas include Davis [Davis83], Kanatani [Kanatani84] and Blake [Blake90], who have all suggested ways of estimating surface orientation from the distribution of orientations of the texture.

b) Topological texture and illumination

The subject of *shape from texture* is not an easy one and it is therefore not surprising that both of the above schools (i.e. both texture gradient and isotropy researchers) have implicitly assumed that the effects of occlusion in topological textures and the effects of illuminant vector variations do not significantly affect image textures. That is they have effectively assumed that image texture results only from surface markings. Exceptions include Kender, Chen & Keller, Choe & Kashyap, Pentland, and Kube & Pentland.

Kender [Kender80] considered surfaces in which texture primitives were either "painted" (parallel to the surface plane) or "pointed" (perpendicular to the surface plane). He did not consider both simultaneously and commented that "textures formed by arbitrary angles to a surface are almost intractable". He did not consider illumination effects on texture.

Chen & Keller [Chen90] state that most shape from texture techniques are based on the assumption that the surfaces are smooth and uniformly covered with flat textured markings or texels. Although Chen & Keller do discuss the use of fractional Brownian motion (FBM) to model the topological texture, and do use a topological model to *generate* test images, they do not use such a model in their shape from texture algorithm directly. Instead they use it to model the intensity texture which could result from either height or albedo variation. They make use of the "*average Holder constant*", which is a fractal-related parameter that changes with scale [Keller87]. This parameter is used to calculate the distance ratio between points on a "planar" surface in order to determine the surface's orientation (i.e. it is a gradient measure). As an intensity model of texture is used it cannot be shown that the average Holder constant is invariant to illumination. This lack of consideration of illumination effects is reinforced by the use of test textures consisting of *computer scaled and rotated* Brodatz images [Brodatz66]. Such rotation ignores illumination effects or at best implicitly assumes that the illumination has been similarly rotated.

Choe and Kashyap [Choe91a] [Choe91b] presented a hybrid *shape from shading*/*shape from texture* technique. They assume that the image is made up of a random texture component and a component due to a smoothed version of the surface (i.e. a surface without topological texture). The smoothed surface is assumed to be Lambertian. An explicit model of the surface's topological texture is not used. Rather, the intensity texture in the surface normal plane is modelled directly as a "fractional differencing model" which has the ability to model anisotropic textures and has a separate variance parameter (see chapter 4).

The key point however, as concerns this thesis, is that Choe & Kashyap effectively assume that a two-dimensional or albedo texture pattern is mapped onto a three dimensional surface and no account is taken of lighting effects. Furthermore, as with Chen & Keller, these assumptions are implicit in the selection of the image test set : images from Brodatz's standard texture album [Brodatz66] were digitised, and *then* subjected to a projection/rotation process. The effects of lighting on three-dimensional or topological textures were therefore not investigated.

Kashyap and his colleagues have also used the fractional differencing model for rotation invariant texture classification [Choe91a] and this will be discussed later in this thesis.

Pentland was the first to report the use of a realistic model of natural topological texture for the purposes of determining shape from texture and texture segmentation. He uses a "spatially isotropic fractal Brownian surface" [Pentland84]. He proposes a proof that the fractal dimension of an imaged texture is identical to that of the components of the surface normals of a spatially isotropic fractal Brownian surface, and goes on to use the fractal dimension as a feature measure for texture segmentation. The two main conclusions of [Pentland84], for shape from shading, are (i) that as real fractal surfaces are fractal over a finite range of scales the perspective gradient of these limits can provide orientation information, and (ii) that fractal dimension can be used as a test for non-isotropy.

In [Pentland86] the use of fractal models for shape from shading is further developed. An image texture measure is presented which is a function of the expectation of the 2nd derivative of the surface normal. This measure is independent of illuminant direction i.e. it is intrinsic to the surface (however it is not clear as to how the illuminant vector is eliminated). As it is affected by foreshortening it can be used to estimate surface tilt and slant. Thus the main conclusion that can be drawn from Pentland's work, for the purposes of this research, is that the fractal dimension of the image of a *spatially isotropic fractal Brownian surface* is identical to that of the components of the surface normals.

Kube and Pentland [Kube88] further investigated the effects of illumination on images of topological texture. They developed a frequency domain model which, given the illuminant vector and the power roll-off factor of a fractal model of the physical texture, allows the two-dimensional power spectrum of the image texture to be predicted. They concluded that the resulting image texture would also be fractal, having a power roll-off factor two less than that of the surface. Their model may be used to predict the directional characteristics of image texture and these predictions have important implications for the majority of texture segmentation and classification schemes (this is discussed further in the last part of this chapter) .

Thus of the shape from texture work Kube and Pentland's fractal-based model would seem to offer the most promising theory. However, before this is described in greater detail the last category of this short review will be presented — i.e. that of scattering theory.

2.1.4. Scattering theory

The effect of rough surfaces on wave scattering has been the subject of many papers and books over the last thirty years. Both electromagnetic and acoustic waves have been investigated and application areas include ultrasonics, sonar, radar imaging, and optics [Ogilvy91]. A vast wealth of literature has been published on this subject. Ogilvy gives an excellent in-depth introduction to this area [Ogilvy91] [Ogilvy87]. Bennett provides a layman's guide to measuring surface roughness of optical and machined components [Bennett89], while Beckmann & Spichino's book [Beckmann63] still provides an often cited reference on the scattering of electromagnetic waves. With many of the titles and abstracts including terms such as "*random rough surfaces*" the area would seem to be extremely relevant to this thesis. However, as the work is concerned with the *scattering* of acoustic or electromagnetic waves, the term "*rough surface*" is defined with respect to the wavelength of the incident irradiation. Thus the typical root mean square (rms) roughness taken into consideration is of the order of 0.2μm or less [Vorburger93], whereas the rms roughnesses of typical test textures used in classification are of the order of millimetres (see [Brodatz66]). The research into scattering is thus concerned with the intimate details of reflection characteristics, whereas for the work described here it is sufficient to assume a reflection characteristic, and use this to investigate the effect of changes in illumination direction on images of comparatively gross surface relief.

Note that some work has been done on composite roughness models [Jackson86] [McDaniel83] in which the surface is modelled as a small-scale roughness superimposed on a higher amplitude, lower frequency, large-scale roughness [Ogilvy91]. The large-scale roughness is normally used to modify the surface normals of the small-scale roughness in a similar manner to Kube and Pentland [Kube88]. Kube and Pentland's theory is however much simpler, as it assumes a Lambertian reflection model — whereas the composite roughness work uses modified normals in the standard Kirchoff or small perturbation theory [Ogilvy91]. The resulting theory is therefore very complex, but it allows the characteristics of the small-scale roughness to be taken into account. Here however, it is mainly the effects that variation in the direction of illuminant incident upon "large-scale" roughness that are of concern. Hence the simpler theory due to Kube and Pentland will be used in this thesis.

2.1.5. Summary

The preceding sections have briefly reviewed four potential areas of image models of topological texture : texture synthesis; texture segmentation and classification; shape from texture; and scattering theory.

Texture synthesis researchers have extensively used three-dimensional models of texture — both for "bump mapping" and generation of fractal landscapes. They have not however generated corresponding models of image texture, which is not surprising given that they are primarily concerned with the *appearance* of their images. On the other hand the texture segmentation and classification researchers might have been more reasonably expected to have developed such models — as "rotation invariant" classification schemes have been reported. However, the majority of this research has not considered problems associated with illuminant variation and surface relief (see chapter 4 for a more detailed review). The third category, shape from texture, yielded Kube and Pentland's frequency domain model which allows the effect of illuminant variation on images of topological texture to be predicted. They assume perfectly diffuse reflection, whereas the last category, scattering theory, is intimately concerned with the details of reflection characteristics. "Surface roughness" in this case refers to variations of the same order as the illuminant wavelength (i.e. hundreds of nanometres). In this thesis however, rms roughness of the test textures is several order of magnitudes higher. In addition the theory is extremely complex. For these reasons it was decided to investigate Kube and Pentland's model in more detail.

2.2. An image model of topological texture

In this section a model of the image of an illuminated fractal surface due to Kube and Pentland [Kube88] is presented. More specifically, an expression for the spectrum of the image that results when such a surface is illuminated by a distant point light source is developed. The theory here differs from [Kube88] in that a simplifying axis-rotation is introduced — this both reduces the complexity of the derivation, and results in an expression for the model, in which the directional effects of lighting are more easily understood. The model is generalised to non-fractal surfaces and this is followed by an examination of the implications that the theory has for texture analysis.

2.2.1. A fractal based image model

A prerequisite for the development of an image model of topological texture is the choice of representation of surface relief. Kube and Pentland chose fractal Brownian motion [Mandelbrot83] to model natural surfaces, as it is widely used in computer graphics [Voss88] [Saupe88]. Their paper essentially applies a simplified version of the Lambertian surface reflectance model to an expression for the power spectral density of the fractal height-map. The theory is split into two parts. *Case 1* considers the situation where the illuminant vector is not perpendicular to the reference plane of the surface texture — allowing the Lambertian reflectance model to be linearised and used in the frequency domain. *Case 2* considers the situation in which the direction of illumination is perpendicular or close to the perpendicular. Here the quadratic term becomes significant and cannot be ignored. The theory becomes complex, involves additional assumptions, and does not yield an expression as a function of either illuminant slant or tilt. Hence only *case 1* will be considered here.

The following theory assumes :

(i) a Lambertian surface (i.e. perfectly diffuse reflection),

(ii) that the fractal Brownian surface $V_H(x,y)$ is band limited such that it is differentiable,

(iii) an orthogonal camera model,

(iv) a constant illuminant vector over the scene, and

(v) a viewer-centred co-ordinate system, in which the reference plane of the surface is perpendicular to the viewing direction.

The following theory first develops a linear model of the intensity image; second, it shows that the two-dimensional partial derivative is a linear operator; third, it introduces a fractal model of the surface; and fourth, it combines the three preceding elements together to provide the frequency domain model.

a) A linear image model of topological texture

The normalised image intensity $I(x,y)$ of the surface is

$$I(x,y) = \mathbf{n} \cdot \mathbf{L}$$

$$= \frac{-p\cos\tau\sin\sigma - q\sin\tau\sin\sigma + \cos\sigma}{\sqrt{p^2 + q^2 + 1}} \qquad (2.1)$$

where

$\mathbf{n}$ = the unit vector normal to the surface at the point (x,y)

$$= \left(\frac{-p}{\sqrt{p^2+q^2+1}}, \frac{-q}{\sqrt{p^2+q^2+1}}, \frac{1}{\sqrt{p^2+q^2+1}} \right)$$

$$p = \frac{\partial V_H}{\partial x} \qquad q = \frac{\partial V_H}{\partial y}$$

$\mathbf{L}$ = $(\cos\tau.\sin\sigma, \sin\tau.\sin\sigma, \cos\sigma)$ is the unit vector towards the light source

τ and σ are the illuminant vector's *tilt* and *slant* angles as defined in figure 2.1.

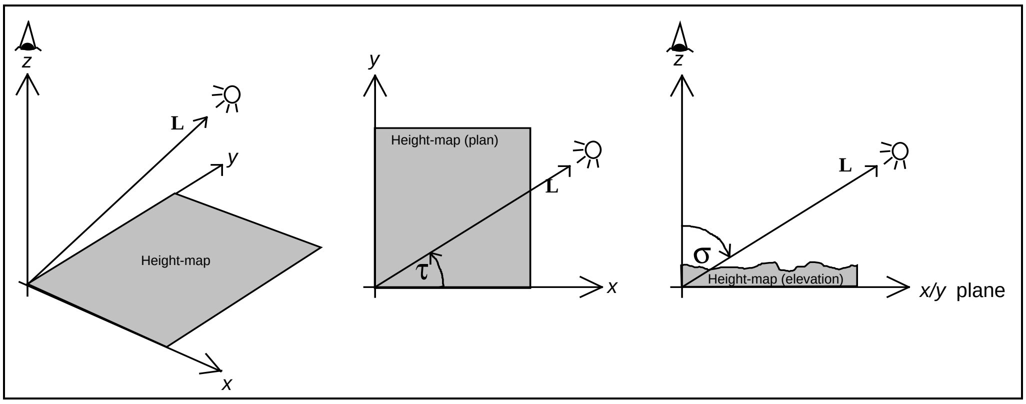

Figure 2.1 - Definition of axis and illumination angles

Now in a departure from [Kube88] and without loss of generality, choose a new axis (*x'*,*y'*,*z*) which is rotated τ about the *z* axis such that the projection of **L** onto the *x*-*y* plane will be parallel to the *x'* axis, as shown in figure 2.2.

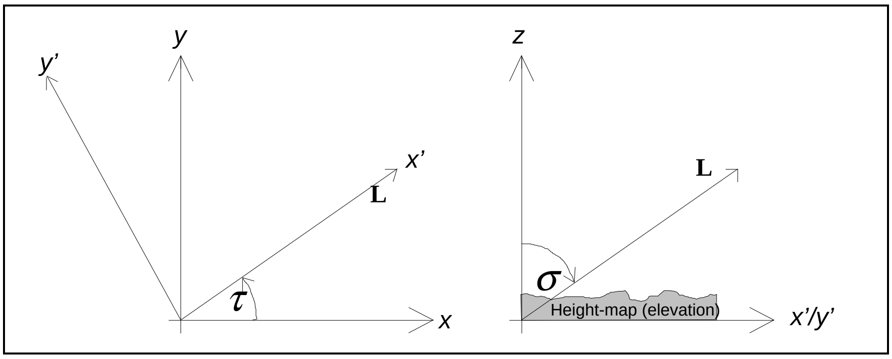

Figure 2.2 - (x,y,z) and (x',y',z) axes.

In this new axis system the expression for intensity simplifies to

$$I(x,y) \;=\; \mathbf{n}\cdot\mathbf{L} \;=\; \frac{-r\sin\sigma + \cos\sigma}{\sqrt{r^2+t^2+1}} \tag{2.2}$$

where

$$r = \frac{\partial V_H}{\partial x'},\ \text{and}\ t = \frac{\partial V_H}{\partial y'}$$

Taking the MacLaurin expansion of $1/\sqrt{r^2+t^2+1}$ yields

$$I(x,y) = (-r\sin\sigma + \cos\sigma)\left[1 - \frac{(r^2+t^2)}{2!} + \frac{9(r^2+t^2)^2}{4!}.......\right] \tag{2.3}$$

A proof of this expansion is provided in appendix A.

Now if the surface slope angles are less than 15°, then $r^2, t^2 << 1$; and the quadratic and higher order terms may be neglected. Note that the error introduced by this approximation, for a slope angle of 15°, is 3.5% (see figure 2.3). With this approximation (2.3) becomes

$$I(x,y) = (-r\sin\sigma + \cos\sigma) \tag{2.4}$$

which is simply the mean, plus a linear contribution of the surface gradient measured in the direction of the illuminant's tilt angle. Thus equation (2.4) is a *linear* model of image intensity, while (2.3) which retains the quadratic and higher order terms is a *non-linear* model of image intensity. It is the former which is of interest here, but both will be referred to in later chapters.

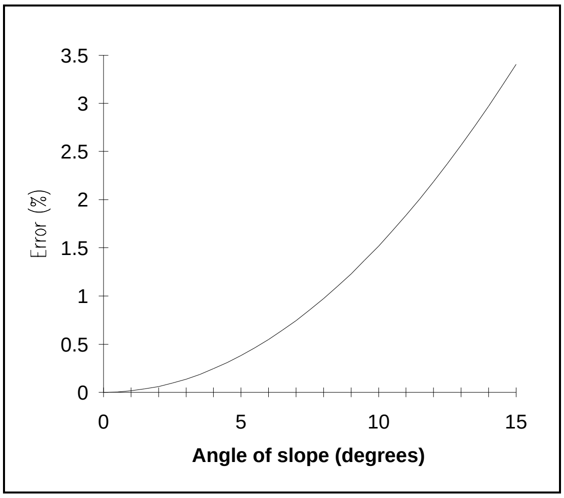

Figure 2.3 - Error due to linear approximation.

Note that if the slant angle is small then sin σ ≈ 0 and the quadratic terms in (2.3) will become important (this is Kube's case 2). For case 1, Kube therefore further assumes sin σ > 0.1, i.e. the illuminant vector **L** is not within 6° of the *z*-axis.

b) The partial derivative operator $\frac{\partial}{\partial x'}$

Consider a single sinusoid surface $V_1(x,y)$ of spatial angular frequency ω_1, angle θ_1 (w.r.t. the *x*-axis), and phase ϕ_1 :

$$V_1(x,y) = \sin[\omega_1(x\cos\theta_1 + y\sin\theta_1) + \phi_1] \tag{2.5}$$

Transforming to the (*x',y',z*) co-ordinate system gives :

$$V_1'(x', y') = \sin[\omega_1(x'\cos(\theta_1 - \tau) + y'\sin(\theta_1 - \tau)) + \phi_1] \tag{2.6}$$

and

$$\frac{\partial V_1'}{\partial x'} = i\omega_1 \cos(\theta_1 - \tau) V_1'(x', y') \tag{2.7}$$

where *i* represents a 90° phase shift.

Taking the Fourier transform yields

$$\mathrm{F}\left[\frac{\partial V_1'}{\partial x'}\right] = i\omega_1 \cos(\theta_1 - \tau) F_1(\omega, \theta) \tag{2.8}$$

where

ω is the angular frequency of the Fourier component

θ is its direction w.r.t. the *x*-axis

$\mathcal{F}[g(x, y)]$ is the two-dimensional Fourier transform of *g(x,y)*, and

$F_1(\omega,\theta) = \mathcal{F}[V_1(x, y)] = \mathcal{F}\left[V_1'(x', y')\right]$

Thus the partial derivative operator $\frac{\partial}{\partial x'}$ is a linear operator, as it does not change either the angular frequency (ω) or the direction (θ) of a two-dimensional sine wave.

c) A fractal model of the surface

From [Kube88] a fractal surface is represented by

$$\mathrm{F}[V_H(x, y)] = F_H(f, \theta) = f^{-\beta_H/2} e^{i\phi} \tag{2.9}$$

where

f is the spatial rotational frequency = $\omega/2\pi$,

ϕ is a random phase element,

β_H is the power roll-off factor[3].

---

[3]Note that for a surface the power roll-off factor β is related to the fractal dimension *D* by : $D = (7 - 2\beta)/2$ [Voss88]. The power roll-off factor will be used in preference to fractal dimension, as

d) Combining the linear intensity model, the partial derivative operator, and the fractal surface model

By superposition, (2.8) and (2.9)

$$\mathrm{F}\left[\frac{\partial V_H}{\partial x'}\right] = i\omega\cos(\theta - \tau)\left(\frac{\omega}{2\pi}\right)^{-\beta_H/2} e^{i\phi} \tag{2.10}$$

Now from (2.4)

$$I(x,y) = -\frac{\partial V_H}{\partial x'}\sin\sigma + \cos\sigma \tag{2.11}$$

Hence if the mean is ignored the Fourier transform of the intensity image is :

$$\begin{aligned} F_I(\omega,\theta) &= \mathrm{F}[I(x,y)] \\ &= \mathrm{F}\left[-\frac{\partial V_H}{\partial x'}\sin\sigma\right] \\ &= -\sin\sigma.\mathrm{F}\left[\frac{\partial V_H}{\partial x'}\right] \\ &= -\sin\sigma.\left[i\omega\cos(\theta-\tau)\left(\frac{\omega}{2\pi}\right)^{-\beta_H/2} e^{i\phi}\right] \\ &= -i\cos(\theta - \tau).\sin\sigma.(2\pi)^{\beta_H/2}.\omega^{-(\beta_H-2)/2} e^{i\phi} \end{aligned} \tag{2.12}$$

The above is mathematically equivalent to Kube and Pentland's case 1, but it contains a simpler expression in terms of θ and τ, which allows the directional effects of illumination to be more easily understood.

Note that the image is predicted to be fractal with a magnitude roll-off of $-^{(\beta_H-2)}/_2$, but that its directional properties have been altered compared with the original surface, i.e. the magnitude of the frequency components is now a function of their angle (θ) in relation to the tilt angle (τ) of the lighting.

e) Generalisation

Although Kube and Pentland used a fractal model of topological texture it is straight-forward to generalise their theory to non-fractal surfaces. If the requirement for the

---

the latter depends upon the measurement method and choice of measurement scale [Voss88] [Mandelbrot85].

surface to have fractal PSD characteristics is relaxed, then as the partial derivative is a linear operator, the Fourier transform of the partial derivative $\frac{\partial}{\partial x'}$ of the surface $V_H\ (x,y)$ is

$$\mathrm{F}\left[\frac{\partial V_H}{\partial x'}\right] = i\omega\cos(\theta-\tau)F_H(\omega,\theta) \qquad (2.13)$$

where

$F_H(\omega,\theta)$ is the Fourier transform of the surface $V_H(x,y)$ which now need not have fractal characteristics.

Hence taking the Fourier transform of (2.11), ignoring the mean, and substituting (2.13) gives :

$$\begin{aligned} F_I(\omega,\theta) &= \mathrm{F}\left[-\frac{\partial V_H}{\partial x'}\sin\sigma\right] \\ &= [-i\omega F_H(\omega,\theta)][\cos(\theta-\tau)][\sin\sigma] \end{aligned} \qquad (2.14)$$

This model (2.14) is now divided into three parts, both to aid understanding, and to facilitate future discussion. The three parts of the model are :

(i) The *surface response component*

$$F_s(\omega,\theta) = -i\omega F_H(\omega,\theta) \qquad (2.15)$$

(ii) The *tilt response component*

$$F_\tau(\omega,\theta) = \cos(\theta-\tau) \qquad (2.16)$$

(iii) The *slant response component*

$$F_\sigma(\omega,\theta) = \sin\sigma \qquad (2.17)$$

Thus Kube and Pentland's model provides theory which allows the influence of illuminant tilt (τ), illuminant slant (σ), and surface characteristics, to be clearly identified.

2.2.2. Implications for texture analysis

In order to aid the design of texture analysis schemes that are robust under lighting variations, it is useful to know which texture characteristics are *intrinsic* to the physical texture, i.e. independent of illuminant, and which are *extrinsic,* i.e. dependent upon the illuminant. In the case of the latter, a knowledge of the behaviour of the texture property under varying illumination conditions, would aid the design of suitable compensation schemes and/or systems that could determine the illuminant's directional characteristics.

The main conclusion of [Kube88] is that a fractal surface with a power spectrum proportional to $f^{-\beta_H}$ produces an image with a power spectrum proportional to $f^{2-\beta_H}$. That is the power roll-off factor (β_I) of an image is predicted to be an *intrinsic* property of a fractal texture — as it is predicted to be independent of the illuminant vector. As far as the directionality of the image is concerned, they merely noted that *"the spectrum depends, as expected, upon the illuminant direction"* and that one of the directional effects could be used for determining the direction of the illuminant. It is however, the directional effects of lighting that are most likely to significantly affect the performance of existing texture analysis schemes — as the majority of texture features surveyed in chapter 4 exploit image texture directionality.

In the following sections the more general model (2.14) is examined with the objective of identifying potentially *intrinsic* or *extrinsic* characteristics of image texture.

(i) The *radial shape* of an image's magnitude spectrum is predicted to be directly related to the radial shape of its surface's spectrum. The term "radial shape" is used here to refer to the shape of a section or slice passing through the coefficient representing the mean. It is purely a function of the surface response F_s, and is therefore an *intrinsic characteristic*. That is, for any value of θ, θ_1 the spectrum in direction θ_1 is

$$F_I(\omega,\theta_1) = k_{\tau\sigma}\left(-i\omega F_H(\omega,\theta_1)\right) \tag{2.18}$$

where

$k_{\tau\sigma}$ is the constant $\cos(\theta_1 - \tau).\sin\sigma$

Thus the radial shape of the log-magnitude/log-frequency graph in any direction θ is predicted to be constant under changes in illumination except for an additive term. This is summarised graphically in figure 2.4.

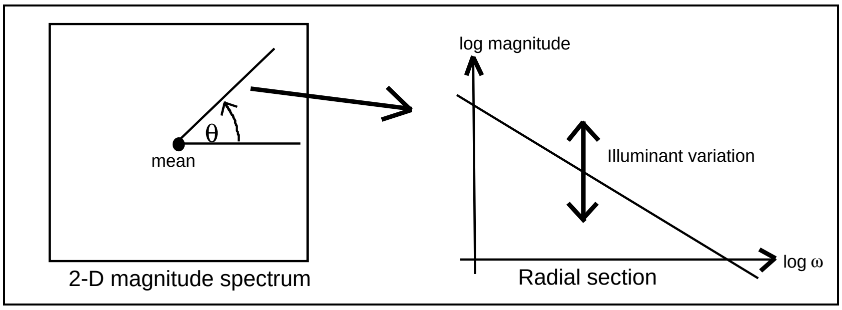

Figure 2.4 -The predicted effect of variation in illuminant direction on the radial shape of the magnitude spectrum.

(ii) The magnitude of any point in the spectrum is a function of illuminant slant (and hence the variance is a function of illuminant slant). So if surface relief and illuminant tilt are held constant the magnitude of a component at any point (ω_1, θ_1) is

$$F_I(\omega_1, \theta_1) = k_{s\tau} \sin \sigma \tag{2.19}$$

where

$k_{s\tau}$ is the constant $-i\omega_1 F_H(\omega_1, \theta_1).\cos(\theta_1 - \tau)$

Thus the absolute values of the magnitude spectrum are a function of σ and any feature based upon these absolute values is an *extrinsic* measure of texture.

(iii) The *angular distribution* of frequency components for an isotropic surface is related to the illuminant tilt angle τ by a cosine function. That is for a "ring" of magnitude spectrum components with radius $\omega = \omega_1$

$$F_I(\omega_1, \theta) = k_{s\sigma} \cos(\theta - \tau) \tag{2.20}$$

where

$k_{s\sigma}$ is the constant $-\omega_1 F_H(\omega_1).\sin \sigma$

In addition the angular distribution of energy of images of anisotropic surfaces (except those surfaces that are perfectly unidirectional) will be a combination of the tilt response and surface directionality[4]. Thus, except for the purely unidirectional case, the directional nature of image texture is predicted to be an *extrinsic*

[4]The term *surface directionality* is used to refer to the angular distribution of a *surface's* variance.

characteristic. This has important implications for texture classification schemes as many make use of directional features.

Normalisation

Some texture classification and segmentation schemes employ normalisation to account for variation in lighting conditions [Greenhill93] [duBuf90] [Laws79] [Weska76] [Haralick73]. Thus they remove any dependence upon absolute magnitude, and so variation in illuminant slant will, according to point (ii) above, be compensated for automatically. However, tilt angle variation cannot be compensated for in the same manner (except if all textures are perfectly unidirectional and all have the same direction). Thus many of texture feature sets that do exploit directionality will *not be* invariant to variation in tilt, unless the test data consists of individually normalised directional textures. This point seems obvious but has not, to the author's knowledge, been considered explicitly in the texture analysis literature (see chapter 4).

2.3. Conclusions

This chapter has briefly reviewed four possible sources of image models of topological texture :

(i) texture synthesis,

(ii) texture segmentation and classification,

(iii) shape from texture, and

(iv) scattering theory.

From these areas a simple frequency domain image model, due to Kube and Pentland, has been selected and presented. This model predicts that image variance and directionality are dependent upon illuminant direction, and only radial shape of magnitude spectra may be intrinsic to the underlying physical texture.

The most important implication that this model has for texture classification and segmentation is that it predicts that many schemes are not invariant to changes in illuminant tilt, and that, unlike slant variation, these effects may not normally be compensated for through the use of normalisation.

However, a number of significant assumptions were made in the derivation of the preceding theory, and validity of the model is therefore the subject of the next chapter.

Chapter 3

An investigation into an image model of topological texture

The previous chapter introduced an image model of topological texture (2.14). This model was used to predict relationships between surface relief, illuminant direction, and image texture — which, if valid, may have significant implications for texture classification. However, the model's derivation relied upon a number of assumptions associated with the projection geometry and the linearisation of the model. In addition shadowing of the surface was ignored. Thus the primary aim of this chapter is to investigate the validity of this model.

For natural textures the most restrictive of the assumptions made, in the author's opinion, are that slope angles are low and shadowing effects are not significant. Hence the investigation reported here paid particular attention to these two aspects.

Chapter 2 divided the model up into three components corresponding to the response of image texture to

(i) changes in surface relief (i.e. changes in topological texture),

(ii) changes in illuminant tilt angle, and

(iii) changes in illuminant slant angle.

In addition to the above responses the model predicts that the radial shape of image magnitude spectra is a characteristic which is intrinsic to the underlying surface relief. That is, it is only a function of (i) and not a function of (ii) or (iii). Hence the objectives of this chapter are to assess the validity of the model by

(a) investigating each of the responses (i) to (iii) above, and

(b) investigating the intrinsic nature of the radial shape characteristic.

Thus this chapter is organised as follows. First, the response of image texture to changes in surface relief is examined. That is the relationship between the magnitude spectra of surface relief and the magnitude spectra of image texture is investigated. Second, the response of image texture to variation in illuminant tilt and slant is presented, and the intrinsic nature of the radial shape characteristic is also examined here. Third and last, the conclusions of the chapter are presented and the implications for texture classification and segmentation re-examined.

3.1. The response of image texture to changes in surface relief

Chapter two's model of image texture (2.14) predicts that the radial shape of image magnitude spectra are determined solely by surface relief characteristics, and therefore may be an intrinsic characteristic of texture. Hence this section focuses upon this important relationship. It was investigated by synthesising height-maps of textures of varying spectra, simulating illumination, and examining the spectra of the resulting images. Physical experiments reported in later sections were used for the investigations into illuminant tilt and slant responses.

As fractal Brownian motion [Mandelbrot85] was used in the development of the image model [Kube88] it is also used here to model and synthesise surface relief. It has the advantages that it is easy to generate and provides natural looking images [Voss88] [Saupe88]. Compare for instance, figure 3.14 with figure 3.23.

The power spectrum of a two-dimensional fractal is of the form $f^{-\beta}$ where β is the power roll-off factor [Kube88], i.e. the log-log PSD plot is a straight line of a gradient of -β. Equation (2.12) implies that, for the fractal case, the power roll-off factor of the image texture (β_I) is related to the topological texture's power roll-off factor (β_H) by

$$\beta_I = \beta_H - 2 \tag{3.1}$$

and this is indeed the main conclusion of [Kube88]. Thus the investigation into the radial shape of texture spectra was restricted to the linear roll-off case. Initially the β relationship (3.1) was examined for a range of surface roll-off factors. The major concern however, was the effect of high slope angles and shadowing. Hence the second and third parts of the experiment investigated these aspects.

However, before the results are presented it is appropriate to describe the process through which they were created.

3.1.1. Image generation

An overview of the process used to generate all of the simulation results described in this chapter is given in figure 3.1.

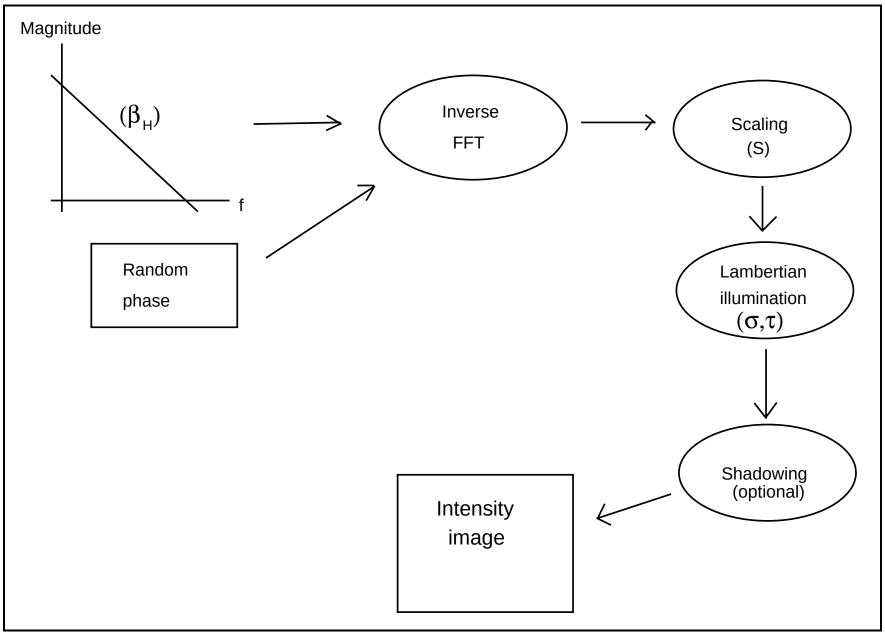

Figure 3.1 - Simulation process showing the major parameters (β_H, S, σ, τ)

All surfaces were synthesised using Fourier filtering [Linnett91a] [Saupe88]. A two-dimensional complex frequency spectrum was created with the desired isotropic power roll-off factor β_H and random phase. This was processed with an inverse Fast Fourier Transform (inverse FFT), and the resulting data where treated as a height-map for input into a Lambertian illumination program followed (optionally) by shadowing.

Details of the illumination program are as follows. The illumination vector was a constant over the scene. Orthogonal projection was used with the viewing direction parallel to the *z*-axis (as in the previous theory). A Lambertian shading model [Newman79] [Rogers85] was employed; the shading equation being derived from (2.1). Estimates of the surface normals ($\hat{\mathbf{n}}$) were calculated from local 2x2 neighbourhoods of height samples as defined below:

$$\hat{\mathbf{n}} = \left(\hat{n}_x, \hat{n}_y, \hat{n}_z\right) \tag{3.2}$$

where

$$\hat{n}_x = V_H(x, y) - V_H(x+1, y)$$
$$\hat{n}_y = V_H(x, y) - V_H(x, y+1)$$
$$\hat{n}_z = 1$$

Multiple reflections were not considered.

The parameters varied in the simulations presented in this chapter are :

β_H - roll-off factor of the log-log PSD of the surface (default 3.5),

S - height scaling factor – used to vary the surfaces' variance, and hence average estimated slope angle (default S = 1),

σ - slant angle of the illumination (default 50°), and

τ - tilt angle of the illumination (default 0°).

3.1.2. The power roll-off factor

The β relationship (3.2) relates power roll-off factors of topological and image textures. Implicit in this relationship is the assumption that the radial shape of the log-log PSD is a straight line, and that the gradient of this line (β) is an intrinsic characteristic of texture. This section reports an investigation into the β relationship (3.1) itself. The intrinsic nature of the PSD's radial shape is further investigated in following sections on slant and tilt angle responses.

A set of simulations was performed where only β_H was varied, the height scaling factor was kept at S = 1 (in order to reduce the effects of the non-linear terms), shadowing was not employed, and the lighting direction was kept constant at τ = 0° and σ = 50°. The power roll-off (β_I) of the resulting images was measured and the relationship between the two parameters estimated. Figure 3.2 shows three of the surfaces displayed as height-maps (where intensity represents height) and the corresponding intensity images. Mean radial sections of the two-dimensional magnitude spectra of the these height-maps and intensity images were obtained by averaging radial sections from θ = 0° to 180°. The resulting plots, together with least square estimates of power roll-off factors, are shown in figures 3.3 and 3.4.

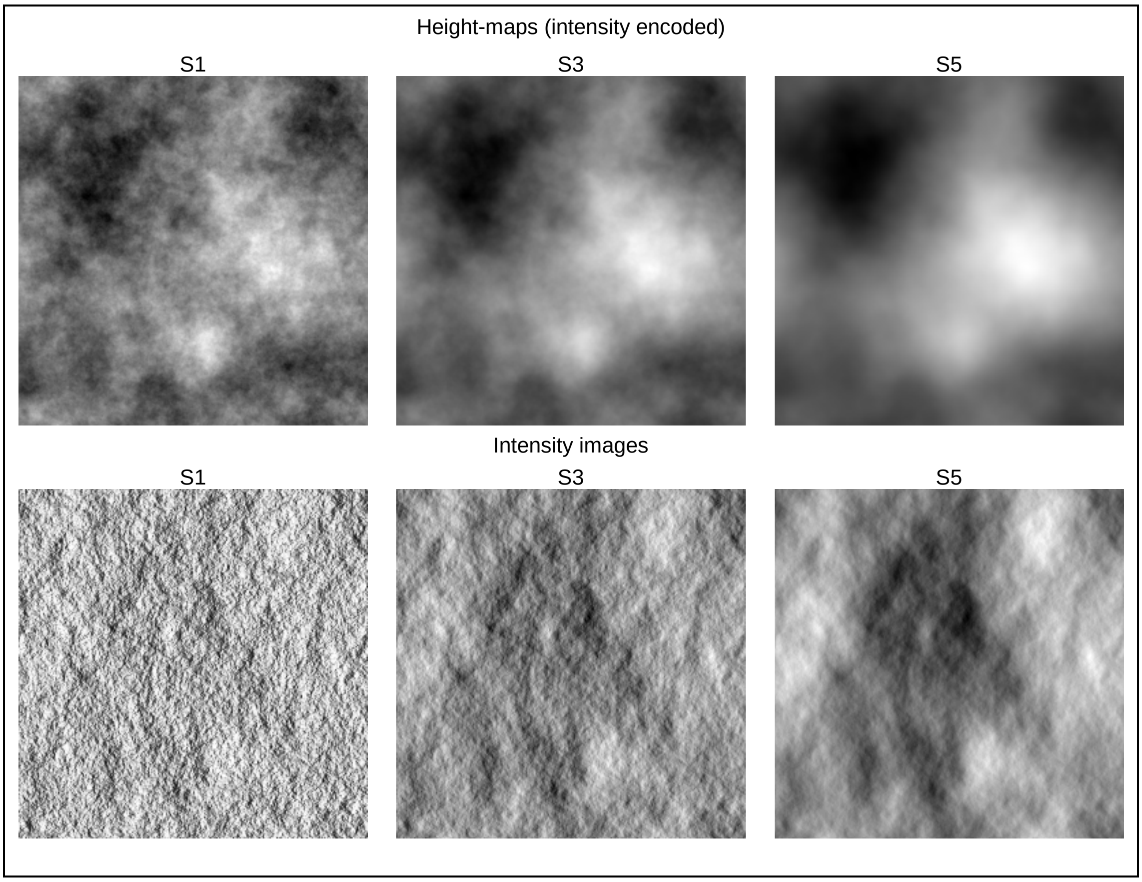

Figure 3.2 - Height-maps $V_H(x,y)$ of the surfaces, and their corresponding synthetically generated intensity images $I(x,y)$. The illumination source is to the right of the surfaces.

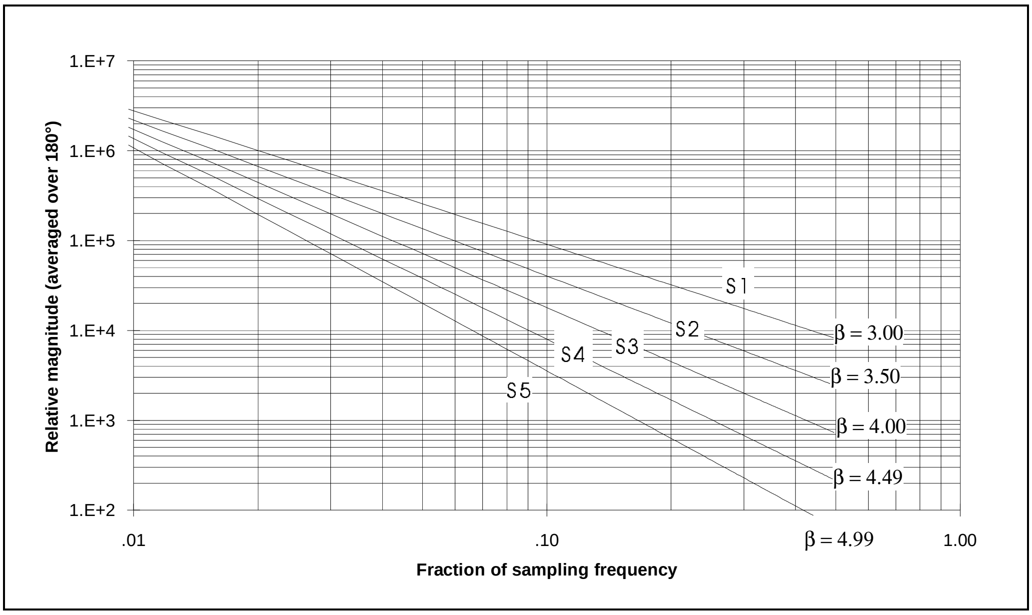

Figure 3.3 - Average radial sections of surface magnitude spectra shown with estimates of β_H.

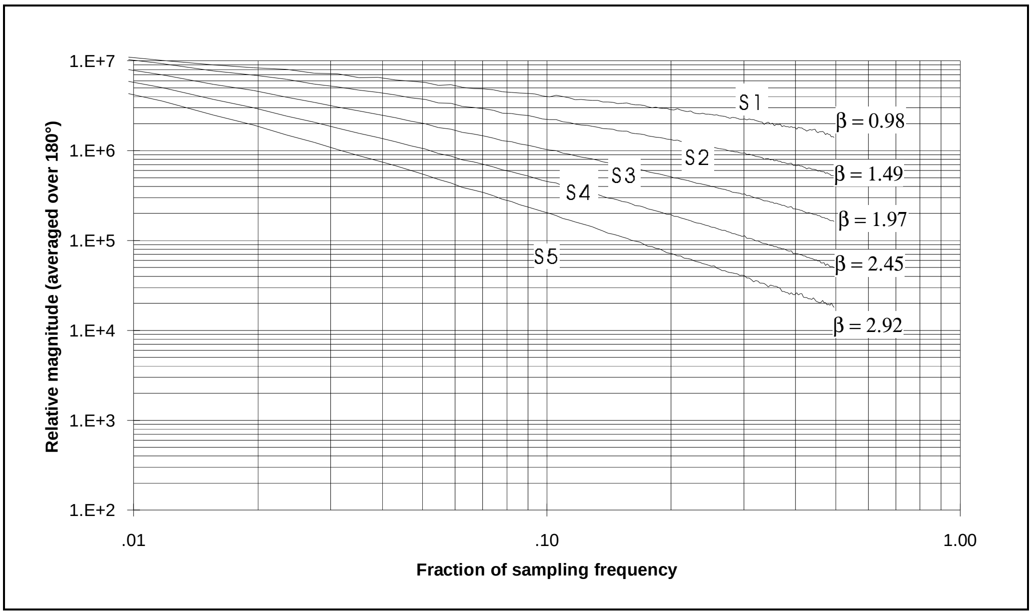

Figure 3.4- Average radial sections of intensity magnitude spectra shown with estimates of β_I.

The above show that both surface and image spectra have linear roll-off characteristics and that, as the theory predicts, the power roll-off factors of the images are approximately two less than their surfaces. Figure 3.5 shows this linear relationship more clearly. Here estimates of power roll-off factor of the images have been plotted against estimates of the original surfaces.

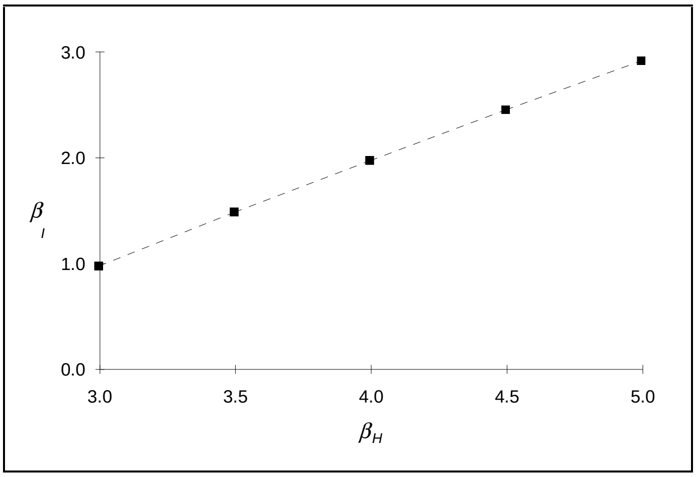

Figure 3.5 - β_I vs. β_H of surfaces S1 - S5.

The least squares estimate of a linear relationship between β_I and β_H (i.e. the best fit straight line to the graph shown in figure 3.5) is :

$$\beta_I = 0.97\beta_H - 1.92 \tag{3.3}$$

This compares favourably with the β relationship (3.1) derived from the model. Note that the *average estimated slope angle* ($\hat{\alpha}$)[1] varied between 2.5° for surface S5, up to 19.0° for surface S1. Thus the relationship between power roll-off factors of topological and image textures has been verified for low to moderate slope angles and no shadowing. The next two sections investigate each of these restrictions in turn.

3.1.3. Large slope angles

The use of the linear image model (2.14) presumes low slope angles and hence low height variance. In order to investigate the effect of larger slope angles the experiment reported above was repeated for increased surface variances. Height-map elements were multiplied by a height scaling factor (S) in the range 1 to 100, while β_H was kept constant at 3.5. The surfaces' *average estimated slope angles* ($\hat{\alpha}$), and *height variances* (s_H^2), are given in table 3.1. The *average estimated slope angles* ($\hat{\alpha}$) are calculated from the angles of the gradients between immediately neighbouring height samples in both x and y directions and averaged over the whole height-map (as defined below).

$$\hat{\alpha} = \frac{1}{n}\sum \frac{|\hat{\alpha}_x| + |\hat{\alpha}_y|}{2} \tag{3.4}$$

where

$$\hat{\alpha}_x = \tan^{-1}[V_H(x,y) - V_H(x+1,y)]$$
$$\hat{\alpha}_y = \tan^{-1}[V_H(x,y) - V_H(x,y+1)]$$

the summation is calculated over the depth-map, and

n is the number of samples contained in the depth-map.

The height variance is defined as :

$$s_H^2 = \frac{1}{n}\sum \left(\overline{V_H} - V_H(x,y)\right)^2 \tag{3.5}$$

where

$\overline{V_H}$ is the mean height of the surface.

[1]The average estimated slope angle ($\hat{\alpha}$) is defined in section 3.1.3

Height scaling factor (S)	1	2	4	10	20	40	100
Surface variance (s_H^2) in distance units^2	7	28	112	700.8	2,803	11,212	70,076
Average estimated slope angle $(\hat{\alpha})$	8.6°	16.2°	28.0°	47.5°	61.1°	71.5°	80.4°

Table 3.1. Average estimated slope angles and height variances, for surfaces with a range of height scaling factors (S)

Figure 3.6 shows sections through surfaces of different height scaling factors. Note that the surface with a height scaling factor of S = 100 has an average estimated slope angle of 80.4° and is therefore not typical of natural surface relief. Nevertheless it is still of value to investigate such extreme data, as they often exaggerate characteristics that might otherwise be overlooked.

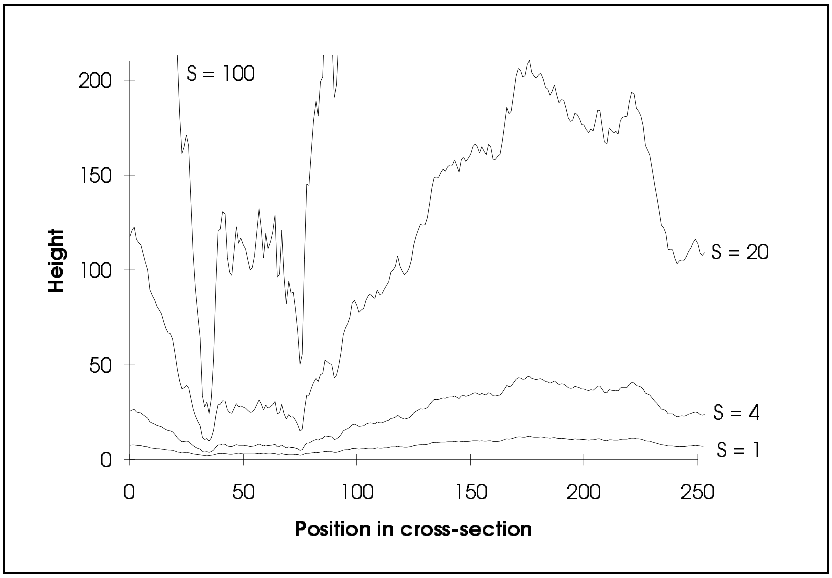

Figure 3.6 - Sections through four surfaces with height scaling factors S = 1, 4, 20 and 100 (note only part of S = 100 is shown for reasons of space)

Each of the surfaces listed in table 3.1 was used as input to the synthetic illumination process. Frequency spectra of the resulting intensity images are depicted in figure 3.7. They show that the gross shape is maintained, but that as the variance of the surface increases that of the corresponding intensity images saturates at a height scaling factor of S = 20. However, the image model (2.14) predicts that image variance is linearly related to surface variance. Not surprisingly, repeating the simulation with the linear illumination scheme implied by the image model — i.e. using equation (2.4) — does not show this saturation effect. Hence it must be due to the quadratic and higher order terms of the Lambertian model (2.3) which are neglected in the linear image model (2.14).

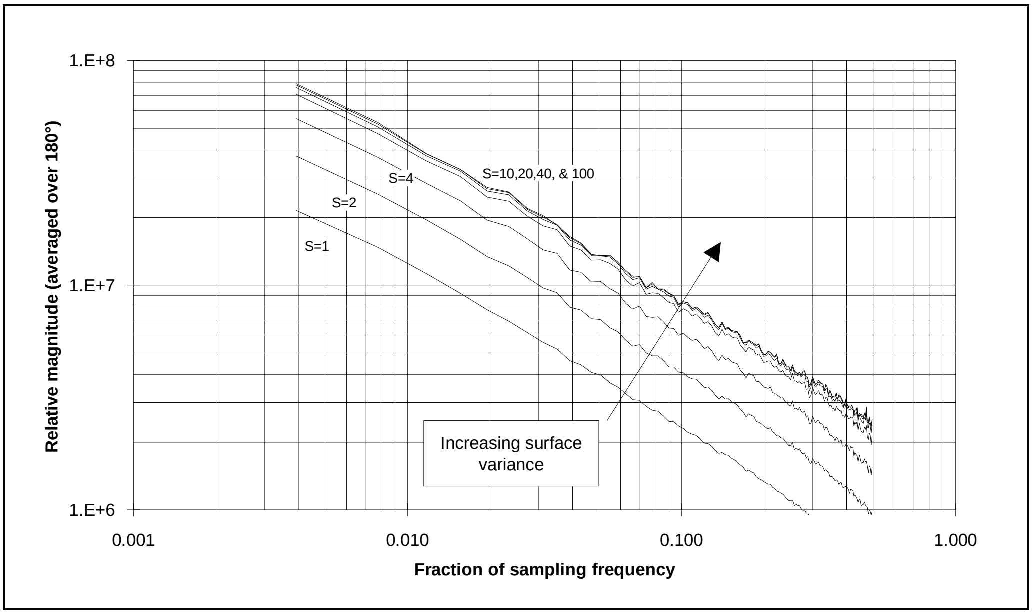

Figure 3.7 - Magnitude spectra of intensity images showing the effect of increasing surface variance.

This supposition is supported by figure 3.8. It shows the effect of increasing the amplitude of a sinusoidal corrugated surface on images generated using (a) the Lambertian model (2.3) and (b) the linear model (2.4). Clearly, as the magnitude of the surface is increased the energy of the intensity radiated by the Lambertian model is reduced compared with its linear companion.

Despite this saturation effect the gross radial shape remains constant over a wide range of average estimated slope angles.

As before roll-off parameters of the intensity images (β_I) were estimated and plotted against surface roll-off factor (β_H) to illustrate the β relationship at a variety of height scaling factors. For clarity the β relationships for only three height scaling factors are shown in figure 3.9. They show that although some deviation from the original relationship is introduced, it is surprisingly small.

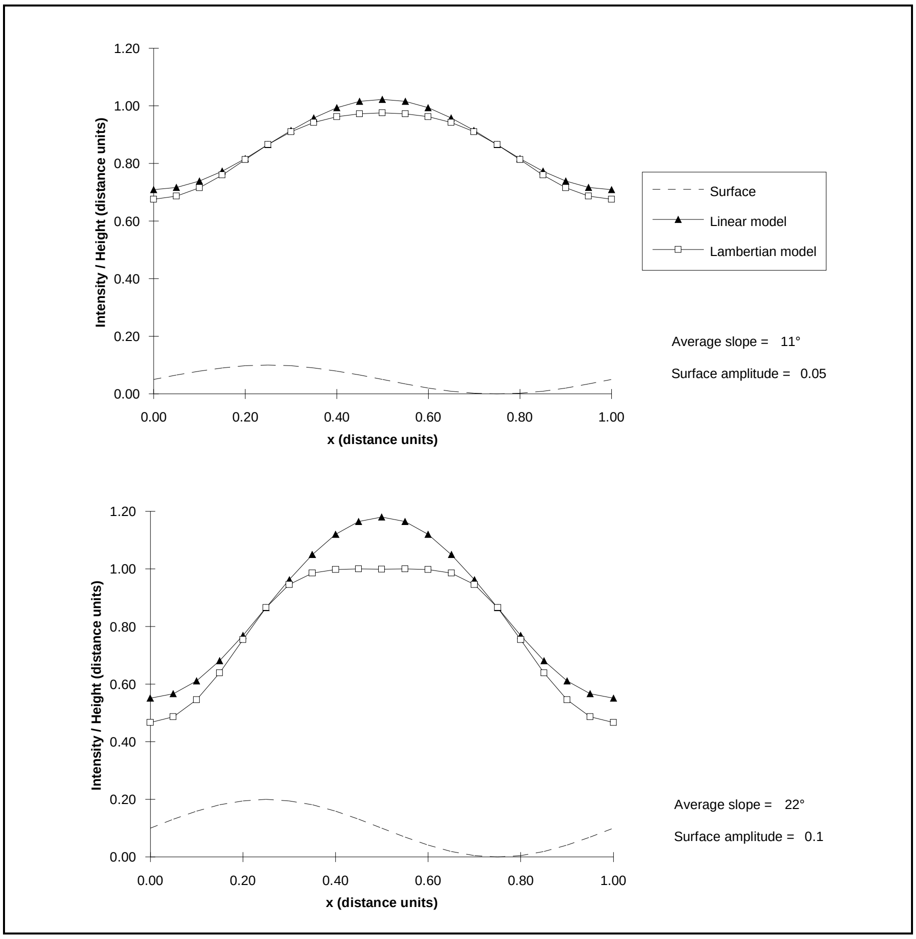

Figure 3.8 - The effect of increasing surface amplitude (from 0.05 to 0.10) on the intensity predicted by Lambertian and linear illumination models

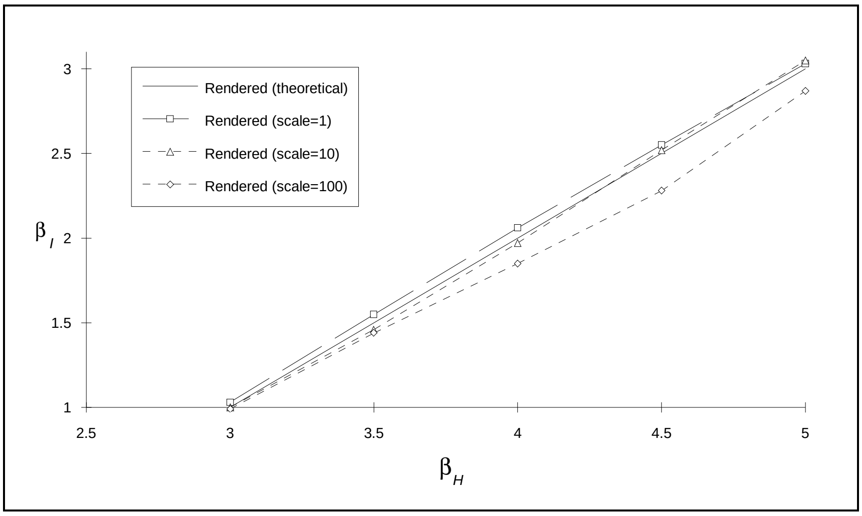

Figure 3.9- The β relationship at height scaling factors S = 1, 10, and 100.

3.1.4. Shadowing

After verifying the β relationship over a range of average estimated slope angles the experiments were repeated with the addition of shadowing. Shadowing was simulated by setting the intensity corresponding to a shadowed height-map sample to zero — thus no account was taken of multiple reflections. The two figures below show "shadowed" intensity images and their spectra for surfaces of constant power roll-off factor but varying height-map variances.

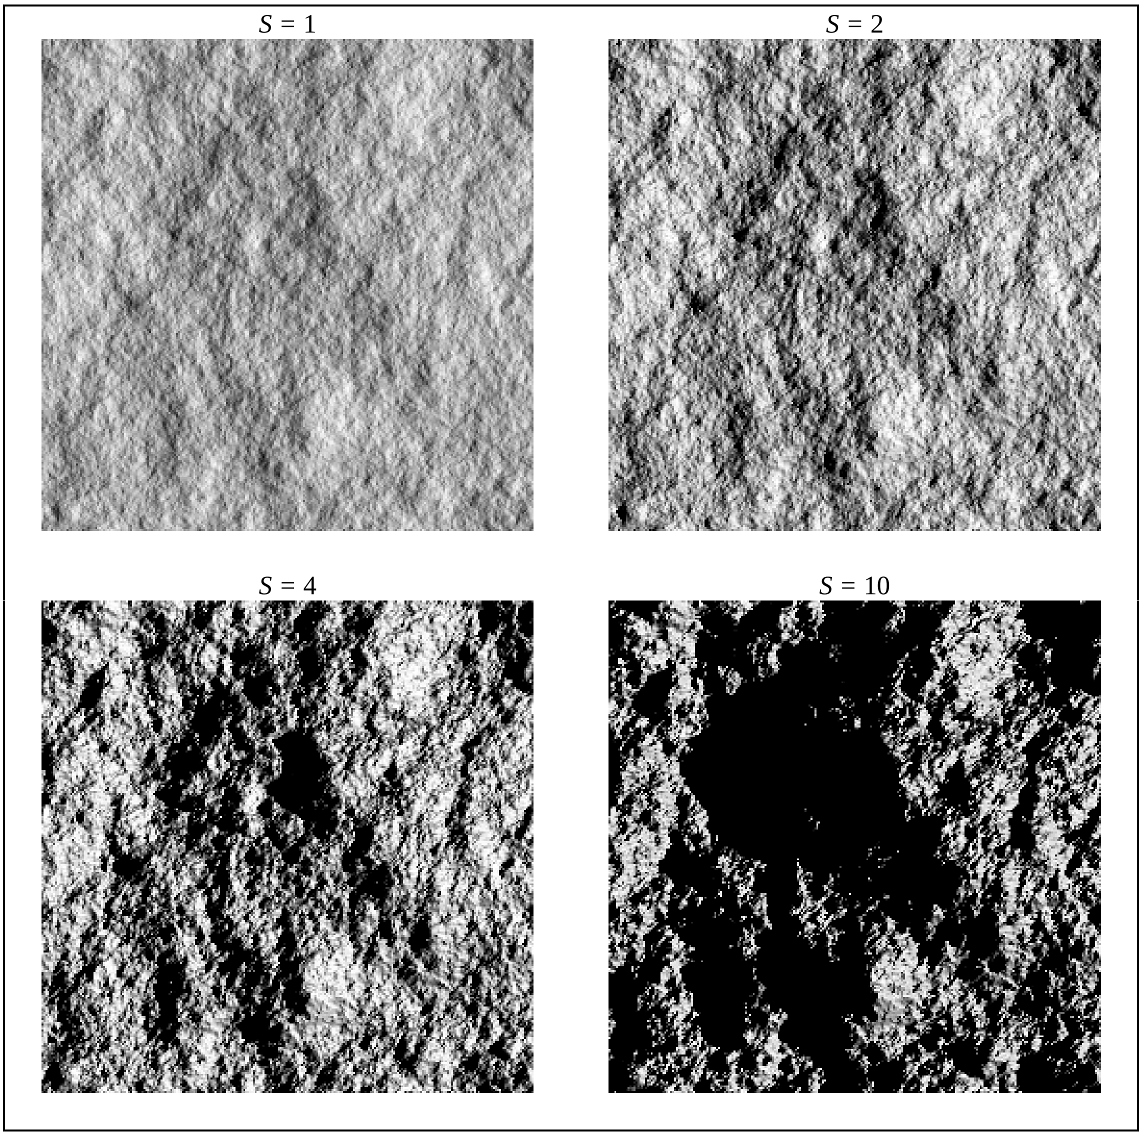

Figure 3.10 - Intensity images (with shadowing) of surfaces of varying height scaling factors (S=1, 2, 4, 10)

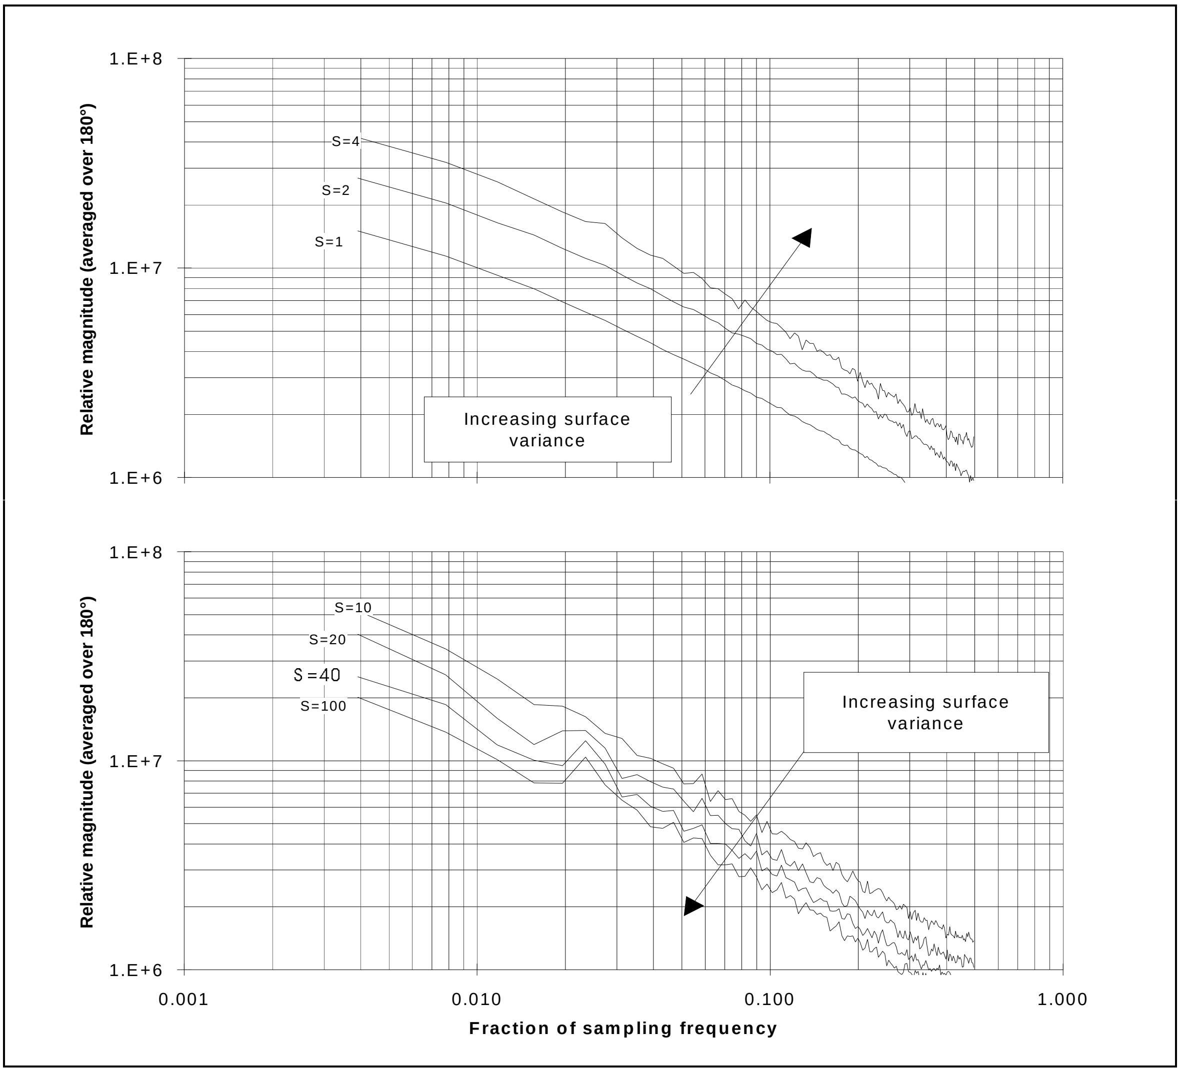

Figure 3.11 - Magnitude spectra of intensity images (with shadowing) showing the effects of increasing the variance of the surface (S = 1, 2, ...40, 100)

The magnitude spectra above illustrate that after a certain point (*S*=10 in the above case) the power spectral density of the intensity images actually decreases as the surface power is increased. However, as before, the straight line nature of the radial shape and its gradient remain largely unchanged.

For clarity figure 3.12 shows the β relationship for only three height scaling factors. From this graph it can be seen that the deviation from the predicted response is surprisingly small given that the theory did not take into account shadowing. This is especially so considering the high degree of shadowing that occurs for surfaces of higher variance — over 80% for a height scaling factor of 100.

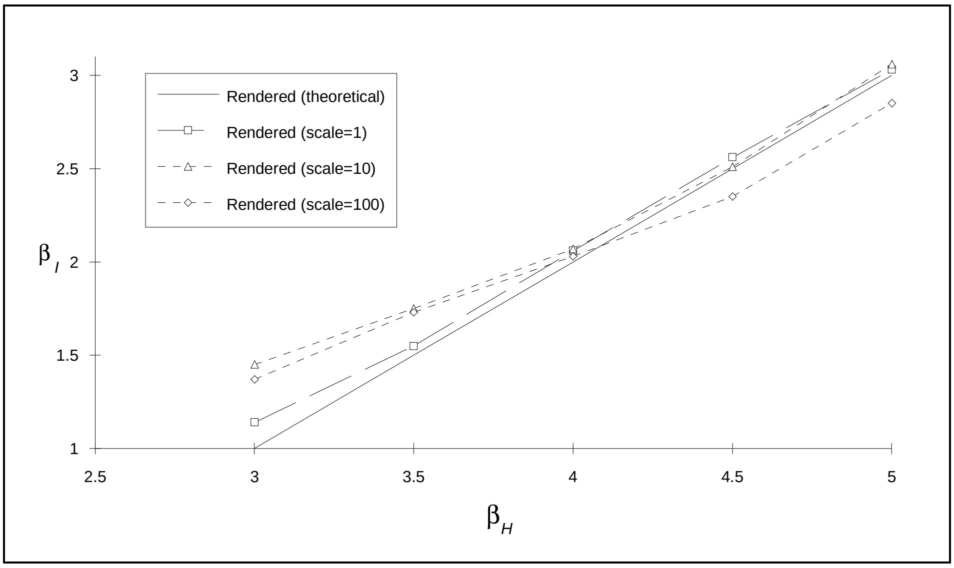

Figure 3.12 - Effect of shadowing, at different surface variances, on the β relationship.

3.1.5. Summary of surface response results

The previous sections have examined the relationship between characteristics of topological texture and image texture through simulation. Of particular interest was the radial shape of magnitude spectra, as the image model presented in chapter 2 predicts that this property is intrinsic to the surface, i.e. it is not affected by variation in illuminant vector. In the case where the radial shapes are straight lines with a roll-off factor β, the relationship reduces to (3.1) :

$$\beta_I = \beta_H - 2$$

It is this β relationship which was the subject of the investigations. Of particular concern was the effect of shadowing and high slope angles. The former had not been considered in chapter 2, while the latter was specifically precluded in order that a linear approximation could be used. The results presented have shown that high average estimated slope angles $\hat{\alpha}$ and shadowing make the simulation output deviate from that predicted by the linear model (2.14). The magnitude of the spectra saturated due to the inclusion of non-linear terms and even reduced when shadowing was included. However, the important result is that the gross radial shape of the spectra was not affected even for high $\hat{\alpha}$.

Thus these results show that for the simulation the β relationship is representative over a wide range of surface variance, but they do not however directly support the suggestion that radial shape is an intrinsic characteristic of texture. The next two sections investigate the effect of variation of illuminant vector, and therefore allow such a proposal to be investigated.

3.2. The tilt angle response of image texture

The preceding section investigated the first part of the image model of topological texture (2.14), i.e. it investigated the surface response component (2.15). This section investigates the validity of the second part of the model — the tilt response component (equation 2.16) :

$$F_\tau(\omega, \theta) = \cos(\theta - \tau)$$

This predicts that the frequency components of a texture, in the same direction (θ) as the tilt angle of the illumination (τ), will be accentuated compared with those components at right angles to this illumination. Thus it implies that an image forming process using directed illumination acts as a *directional filter* of texture. Such an effect is likely to have important implications for texture classification schemes. It implies that the directional properties of image texture are not intrinsic to the surface, but that they are considerably affected by variation in illuminant tilt. This is unfortunate, as the majority of the feature sets reviewed in chapter 4 exploit directional characteristics.

Unlike the previous study, which was primarily required to vary surface relief in a controlled manner, the main requirement of this investigation is much simpler — that of varying the illuminant's tilt angle. Thus both simulation and laboratory experiment were used. Simulation was employed as before to selectively examine the effect of the non-linear terms and shadowing. Physical experiments were conducted to provide confidence that the simulations were reasonably representative of the behaviour of real texture, and to investigate a number of differing surface reliefs.

As with the discussion of the surface response, the tilt angle response is first investigated for low slope angles followed by an investigation into the effects of increasing the surface variance and the addition of shadowing.

3.2.1. Low slope angles

An isotropic surface was generated and illuminated synthetically as before (S = 1, β = 3.5, σ = 50° & τ = 0°). Figure 3.13 shows a polar plot of the FFT of the resulting image texture, in which each point on the graph represents the sum of the magnitude coefficients in one direction (i.e. for one value of θ : the angle of the frequency component).

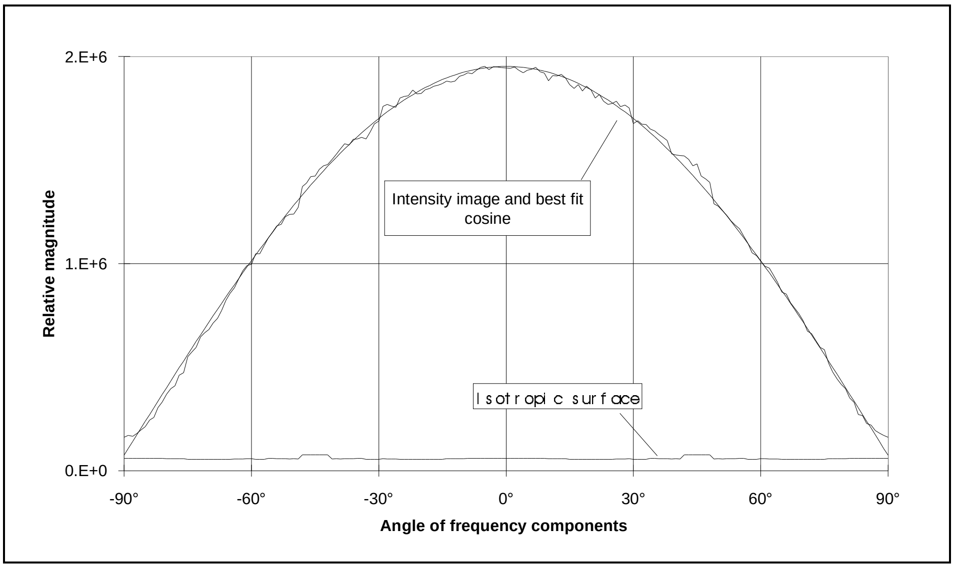

Figure 3.13- Polar frequency plot of image texture (τ= 0°), and corresponding best fit cosine (original surface also shown).

The directionality in the image is clearly evident in the polar plot shown above, especially when the graph is compared to the almost flat plot of the original surface. As predicted by the image model the polar response is greatest in the direction of the illuminant tilt and it follows a cosine distribution very closely. However these data do not illustrate the effect of variation in illuminant tilt angle : figure 3.14 shows images for illuminant tilt angles of 0° and 90°.

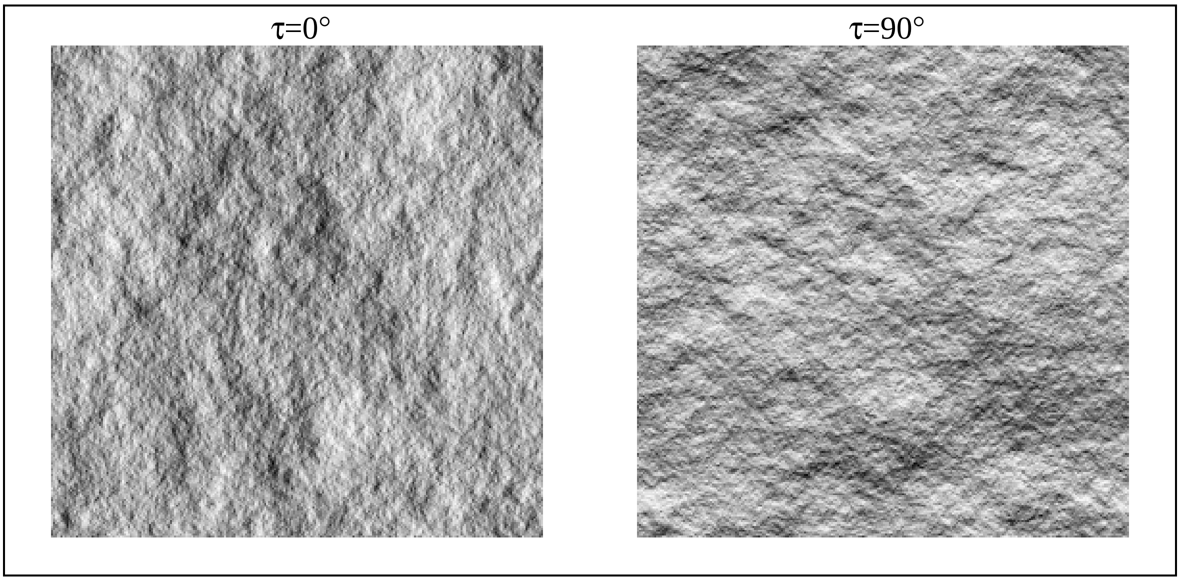

Figure 3.14 - Intensity images showing variation with tilt angle (τ)

The effect on these images could not be described as dramatic but it is clearly discernible. However, in the frequency domain the response to a change in tilt is much more obvious as shown in the polar plots below.

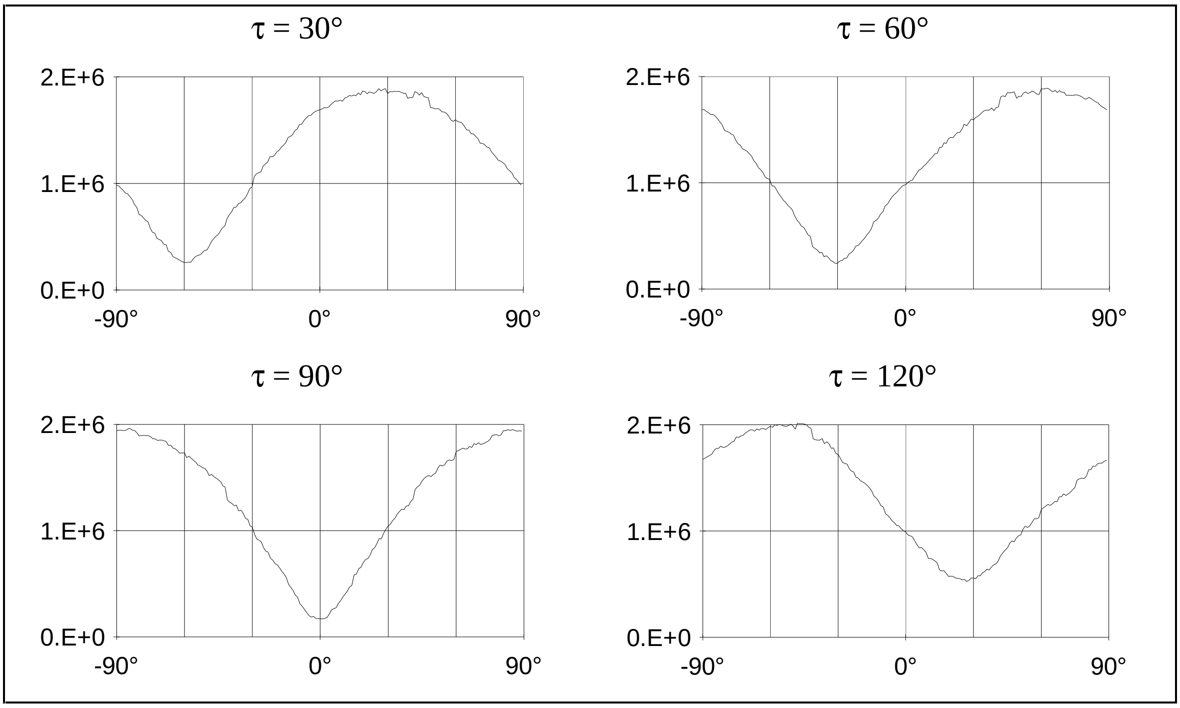

Figure 3.15 - Polar frequency plots of image texture showing the effect of variation in illuminant tilt (τ). Axes are as figure 3.13.

The above demonstrate that, in simulation, the tilt angle responses of images of isotropic topologies closely follow the directional characteristics predicted by the model. These results are not surprising as the synthetic surface had an average estimated slope angle of 8.6° — and the effect of the quadratic and higher order terms neglected in the linear model(2.14) would be small. The next section therefore examines the effects of larger slope angles on the directional characteristics of image texture.

3.2.2. Large slope angles

As before slope angles were increased simply by multiplying the original height-map by a height scaling factor (*S*), which naturally also increases the surface height variance σ_H^2. Figure 3.16 shows polar plots of the two dimensional magnitude spectra of the resulting images.

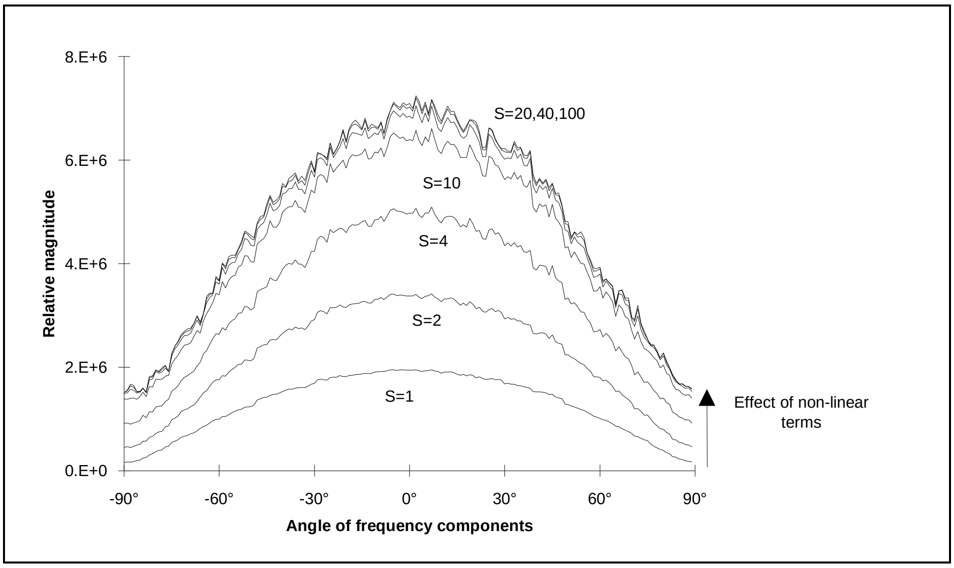

Figure 3.16- The effect of increasing average slope angles on the polar plots of magnitude spectra

These results show that large slope angles affect the cosine form of the image textures' directional characteristic very little. Increasing the surface variance increases image variance as predicted, except that, as was the case in the previous section on the β relationship, the image variance saturates at S = 20. This is due to the non-linear Lambertian illumination expression used in the simulation, as discussed in the previous section. However, the most interesting non-linear directional effects occur at $\theta = \tau \pm 90°$. Here the linear model predicts that *all* components will be filtered out, but figure 3.16 shows this is not the case. Repeating the simulation using the linear illumination model removes the "saturation" and modified directional filtering effects, and shows that both are due to the quadratic and higher order terms of the Lambertian model (2.3).

If a surface consisting solely of components with a direction θ = τ ± 90° is considered, then

$$r = \frac{\partial V_H'}{\partial x'} = 0 \tag{3.6}$$

substituting (3.6) into the Lambertian model (2.3), and ignoring the mean term gives

$$I(x,y) = \cos\sigma\left[-\frac{t^2}{2!} + \frac{9t^4}{4!}\cdots\cdots\right] \tag{3.7}$$

Thus the image will also consist only of components at θ = τ ± 90°. They are generated by the square and higher order t terms that are neglected in the linear model (2.14). For an image of an isotropic surface these terms will give rise to the "non-linear effects" seen at θ = τ ± 90° in figure 3.16, and will naturally become more significant at higher slope angles.

Therefore the "directional filtering" effect is reduced at higher surface variances. When shadowing is included it is further reduced as is shown in the next section.

3.2.3. Shadowing

As in the previous section, on the surface response, shadowing was investigated through the use of simulation. From figure 3.17 it is clear that shadowing only affects the polar plots significantly for height scaling factors of S = 10 and above. The polar plots of surfaces with a height scaling factor of S = 4 or less resemble their non-shadowed counterparts very closely. However, for surfaces with a height scaling factor of 10 and above, the variance of shadowed images actually reduces as the surface variance increases. This echoes the results obtained for the β relationship. Note however, that these polar plots still retain their cosine characteristic, but that they would be better represented by a raised cosine as the minima (at θ = τ ± 90°) increase with surface variance. Thus the "directional filtering effect" is most severe at low slope angles.

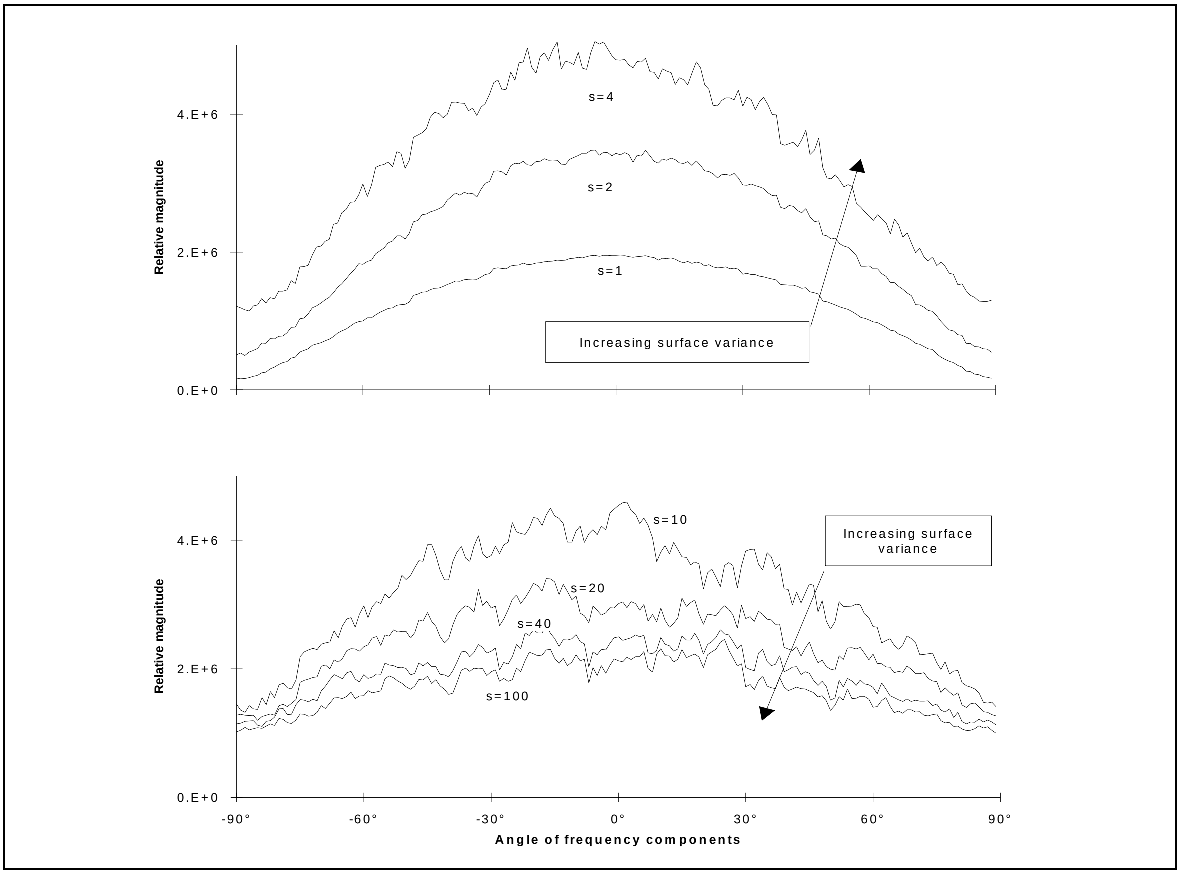

*Figure 3.17- Magnitude vs. angle of frequency components of **shadowed** images for various height scaling factors.*

3.2.4. Four physical textures

All of the results presented so far have been obtained via simulation. Its use has enabled the power roll-off and variance of surface textures to be precisely controlled in order that non-linear effects could be investigated. Shadowing has also been selectively investigated. These experiments would have been either difficult or impossible to perform with real textures. However, the exclusive use of simulation may result in false conclusions being drawn due to the incorrectness of either explicit or implicit assumptions. Hence in this section results of laboratory experiments are presented using four different samples of texture. These samples were selected using the following criteria :

(i) The textures had to be isotropic in appearance to minimise their impact on the directional characteristics of the image textures.

(ii) The "scale" of each texture had be such that it could (a) be detected by the imaging system, and (b) was not so large that a representative sample of it would not fit onto one of the 60 cm square mounting boards used in the experiment.

(iii) The textures had to be of a material that could be spray painted.

(iv) The texture samples had to be "globally" flat.

Images of the four textures are shown below in figure 3.18.

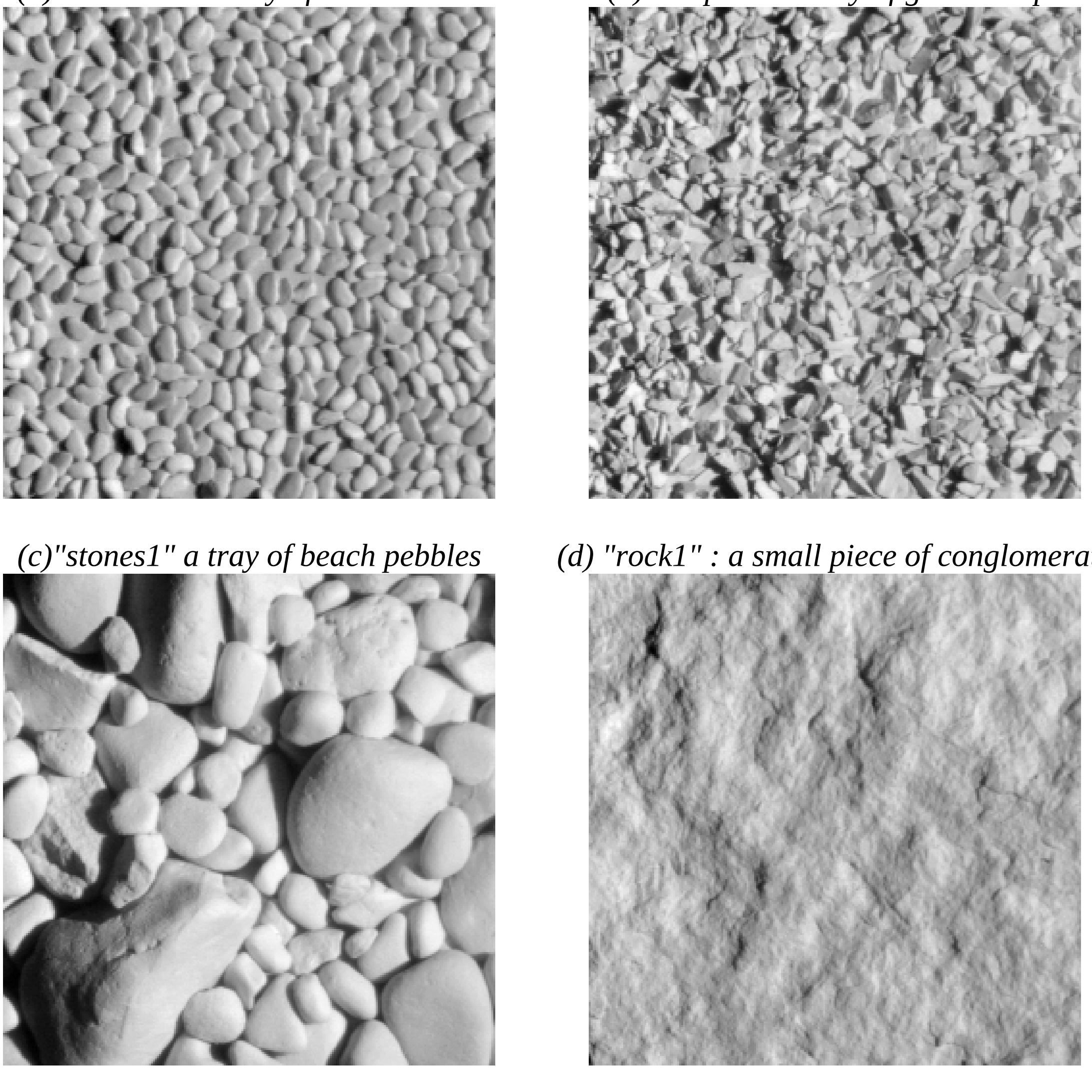

(a) "beans1" : a tray of butter beans

(b) "chips1" : a tray of gravel chips

(c)"stones1" a tray of beach pebbles

(d) "rock1" : a small piece of conglomerate

Figure 3.18 - The test textures.

If experimental results are to be of value then it is important that the phenomena that they exhibit are seen to be due to the process under investigation rather than the experimental procedure or analysis. The following section therefore describes the experimental set-up and the analysis techniques employed after which the results are presented. As directional characteristics are important here special attention was paid to their possible artificial introduction, both in the capture of the images and the ensuing frequency domain analysis.

a) Experimental technique

General set-up

Each of the textures was sprayed matte white to eliminate any albedo texture and to provide an approximately Lambertian reflectance characteristic. Images (512x512x8bit) were captured using a CCD[2] camera with a 40 mm lens (aperture = f11) connected to a frame store mounted in a workstation. The texture samples were mounted as shown in figure 3.19.

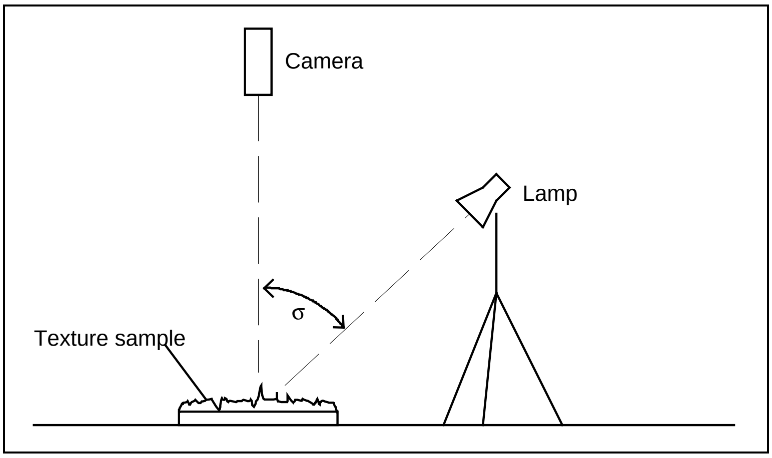

Figure 3.19 - Experimental set-up

That is they were mounted perpendicularly to the camera's line of sight at a distance of 3.3m; and illumination was provided by a 500W lamp, 1.6m from the subject. The position of the illumination was varied in terms of tilt and slant angles, and all other parameters were kept constant.

[2]Charge coupled device

Compensation for illuminant intensity variation

As the illumination source was mounted relatively close to the texture, and, as its lighting pattern was unknown, the variation in intensity of illumination incident on the textures' surfaces was investigated. It was especially important to remove any directional trends — as spectral leakage in the FFT process could *smear* the low frequency components due to the illumination trend to affect higher frequencies : thereby giving the illusion of a general trend over the whole frequency range. Variation in illumination was assessed by taking "registration images" of a flat matte white board;. A variation of 18% was observed in grey-levels. Registration images were therefore captured for each texture image and used to compensate for illumination intensity variation. Each texture image grey-value was divided by the corresponding registration grey-value.[3]

b) Spectral estimation of image textures

The images shown in figure 3.18 are random in nature. Estimation of their spectra therefore becomes the problem of spectral estimation of random fields [Brigham88] [Marple87] [Kay81]. The main criteria for this estimation task are :

(i) directional artefacts should be minimised,

(ii) general trends of the spectra are more important than specific detail,

(iii) *changes* from one magnitude spectrum to another, due to variation in illuminant tilt and slant angles, are more important than the absolute accuracy of the spectra themselves.

Unfortunately the raw application of a two-dimensional FFT routine to the image textures presents two problems : firstly the variance of the coefficients appears high relative to the underlying trend, and secondly large directional artefacts are introduced at $\theta = 0°$ and $90°$.

Directional artefacts

Directional FFT artefacts can be detected simply by rotating a digital image of texture and performing FFTs on the rotated and original images. Their polar plots (normalised with

---

[3]Note registration images were first normalised to a mean of 1.0 — by dividing each registration image pixel by the original registration image mean.

respect to θ) will then be identical except for any directional artefacts introduced. An example for the texture *rock1* for a 45° rotation is shown below.

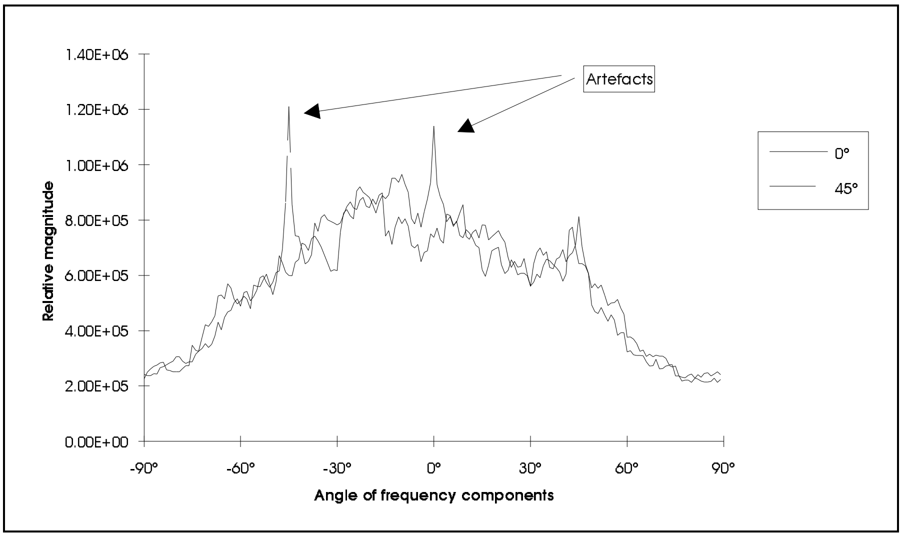

Figure 3.20 - Directional artefacts of raw FFT process

The above artefacts are caused by discontinuities formed by the straight edges of the image, and can be reduced by the application of a circular window [Huang72] [Brigham88, p252]. The next figure shows a sample of the results obtained by applying a circular window to a sequence of images.

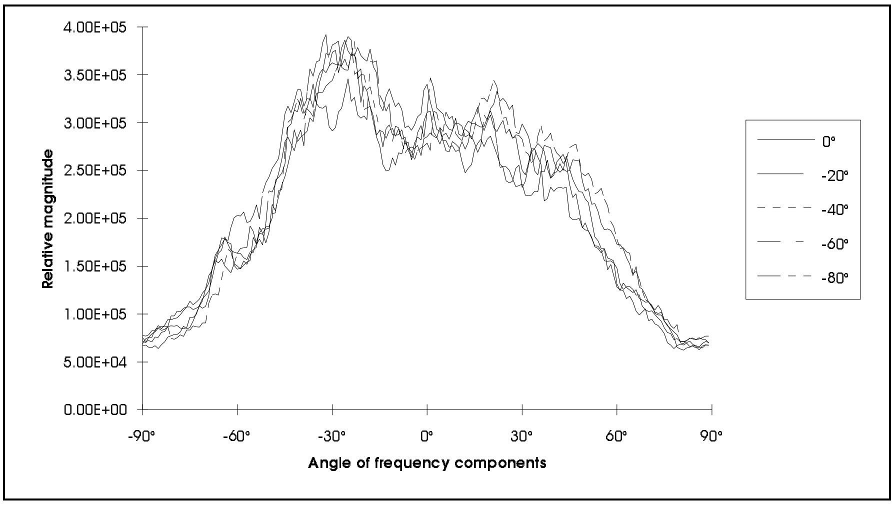

Figure 3.21 - Effect of a circular Hann window

These images were generated by rotating the camera about its viewing axis in 10° steps. Note that each plot has ben off-set by the appropriate camera off-set angle. All of the resulting polar plots were similar to the sample shown. Their similarity shows that no significant directional artefacts are introduced in either data capture or analysis *providing* a circular window is used.

Variance of Fourier coefficients

As the texture images are effectively random fields, it is not surprising that estimates of spectral coefficients obtained via the straight forward application of an FFT routine appear to exhibit high variance relative to the underlying trend. Standard methods of reducing variance of classical periodogram PSD estimators involve either spatial or frequency averaging. The Welch periodogram [Welch67] is straightforward to implement and has proven to be a robust estimator [Marple87]. It divides the data up into segments which overlap each other by 50%. The segments are windowed (using a circular Hann window) to reduce spectral leakage [Marple87], and transformed with an FFT to provide multiple periodograms which are averaged together. The figure below shows the radial sections of spectral estimates using three differing segment sizes. Note that a 512x512 image was used and so "one 512x512 segment" refers to a straight (non-averaged) FFT process. It has been plotted for comparison purposes.

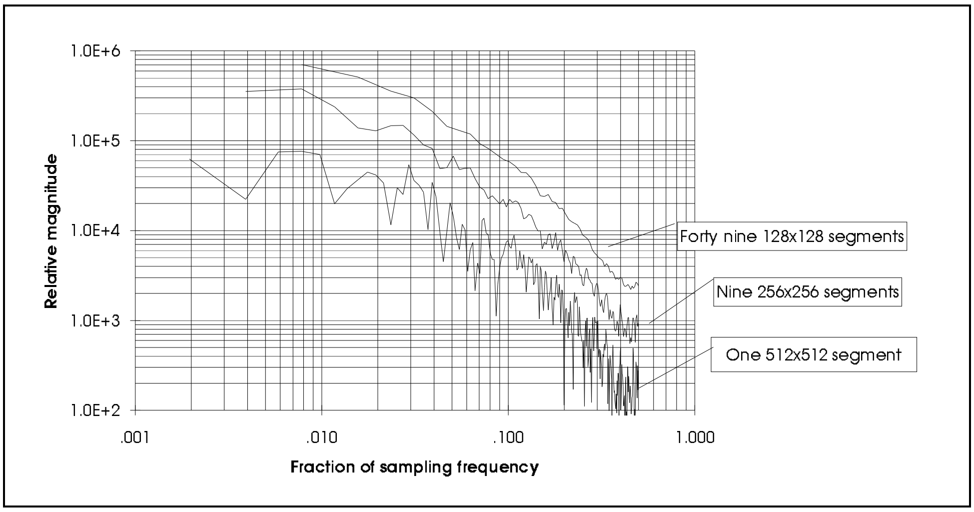

Figure 3.22 - Spatial averaging of magnitude spectra (note the spectra have been displaced vertically for display purposes)

The above figure shows that, as expected, reducing the segment size reduces the variance. In this thesis overall *trends* in response to illuminant variation are of interest rather than detail of spectra. The ability to reduce the "spread" of plots is valuable as it allows differences between graphs to be more easily observed, rather than being obscured by their own variance. Hence the Welch periodogram was used for the generation of all spectral estimates of images of physical textures.

c) Tilt response : experimental results

The illuminant's tilt angle (τ) was varied in 10° steps over 180° for the four textures. The experimental set-up was as described in (a) above. Two examples of the resulting images are shown in figure 3.23. Magnitude spectra of the images were estimated using the Welch periodogram method using forty nine overlapping segments. Examples of the polar plots of the two-dimensional spectra of *rock1*, are shown in figure 3.24.

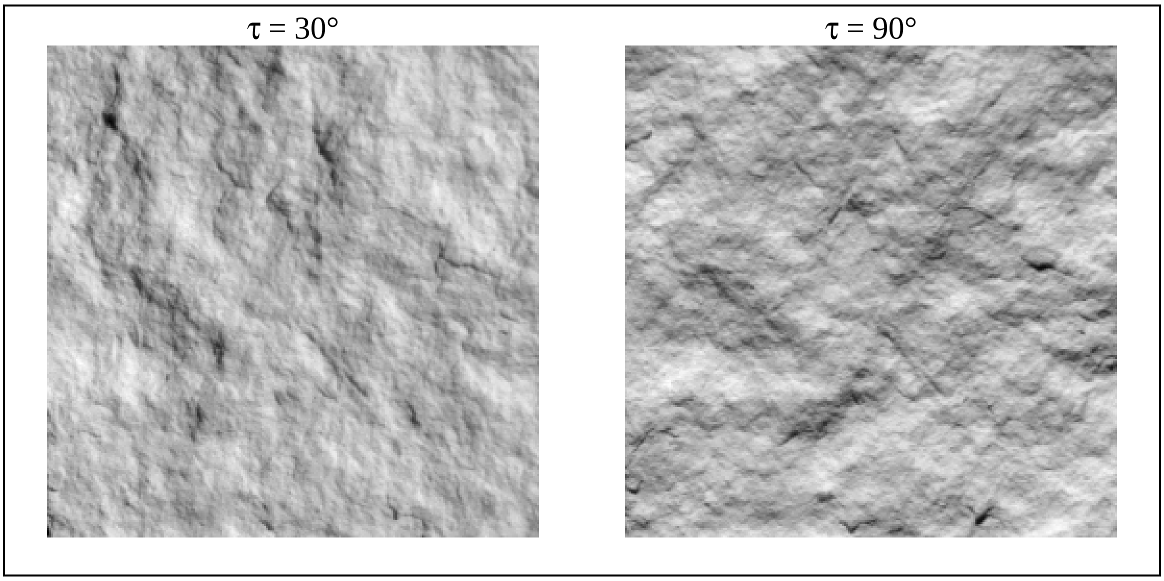

Figure 3.23 - Images of "rock1" captured at two different illuminant tilt angles

As predicted, illuminant tilt clearly has a considerable impact on directionality of image texture *rock1* (note that the angular position of the magnitude peak follows τ). What is perhaps more surprising however, is the similarity of the above plots to those obtained via simulation. Compare, for instance, the τ = 30° plot above with that of figure 3.15; both resemble a raised cosine and both have clear minima within a few degrees of -60°.

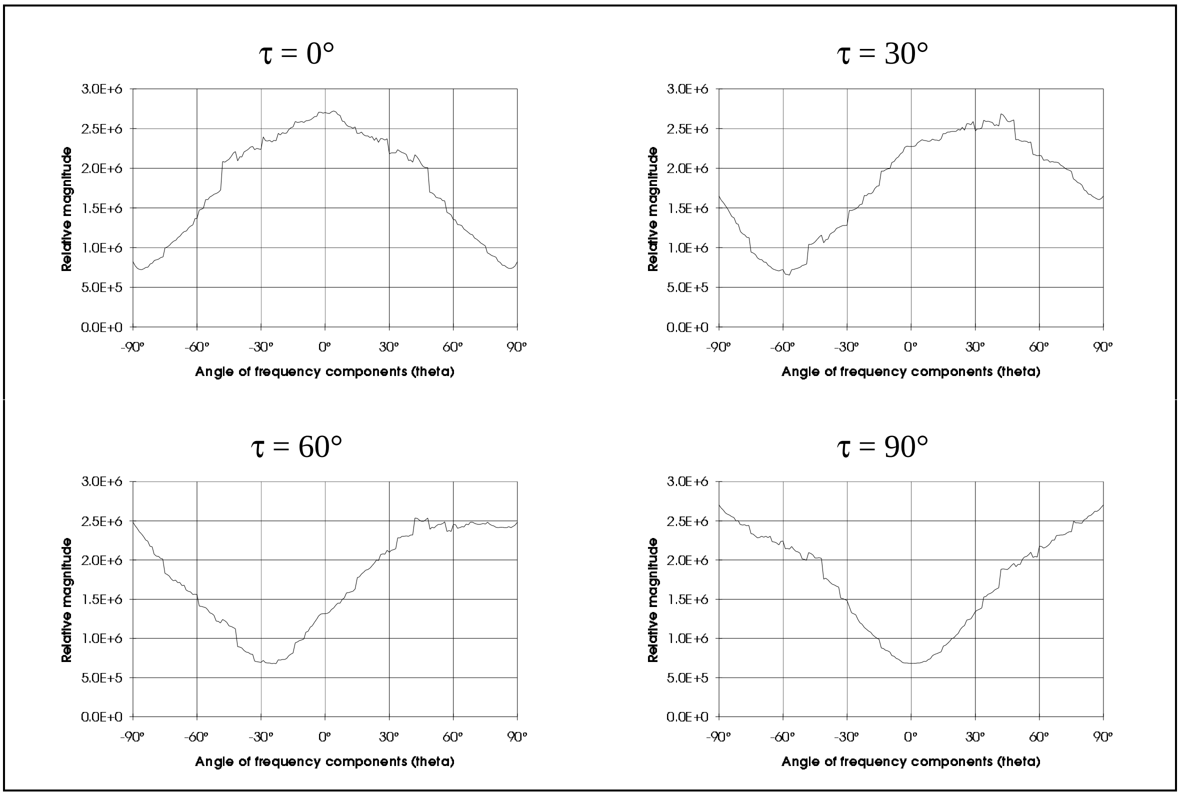

Figure 3.24 - Effect of illuminant tilt angle on image directionality (rock1)

The cosine relationship is more obvious in figure 3.25, in which magnitude has been plotted against cos(θ - τ).

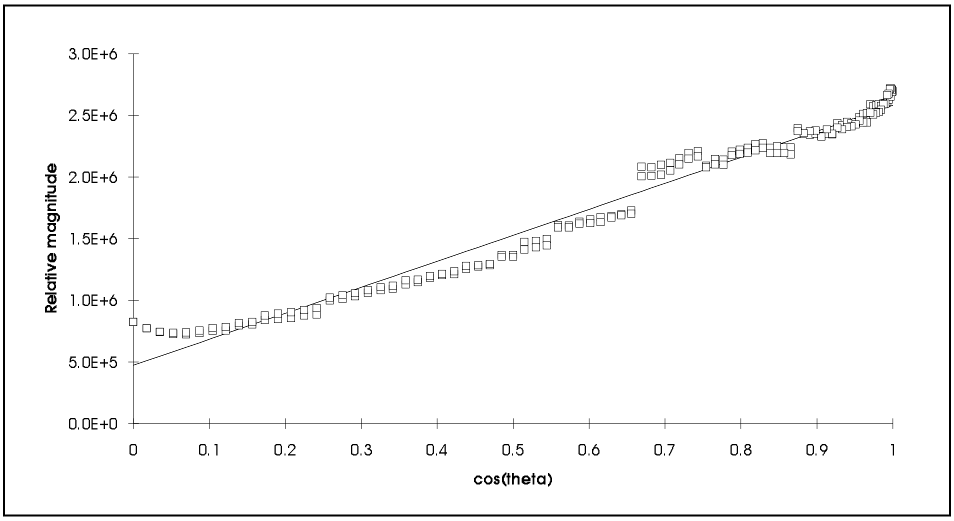

Figure 3.25 - cos(θ - τ) relationship for rock1, and best fit straight line
$y = m_\tau \cos(\theta - \tau) + b_\tau$ *(where y is the relative magnitude)*

It shows that there is an approximately linear relationship between magnitude and cos(θ - τ). Here the magnitude of the "platform" of the *raised* cosine can be determined from the y-intercept of the graph (b_τ). The platform was caused in the simulation by non-linear terms — the height of the platform being related to the *average estimated slope angle* of the original surface. Increasing the slope angles increases the contributions of the non-linear terms and results in a higher platform. Thus if an image of a surface with apparently higher slope angles such as *stones1* were captured, its spectra would be expected to exhibit a higher platform. Table 3.2 below shows the slope and intercept estimates for all four textures including *stones1*. The estimates were obtained using least squares linear regression.

	rock1	beans1	chips1	stones1
slope(m_τ)	2.1E+6	2.2E+6	2.2E+6	1.7E+6
y-intercept(b_τ)	0.47E+6	1.7E+6	3.2E+6	2.7E+6

Table 3.2. Best fit raised cosine parameters for $y = m_\tau \cos(\theta - \tau) + b_\tau$

Figure 3.26 shows polar plots of the four image textures together with their best-fit raised cosines.

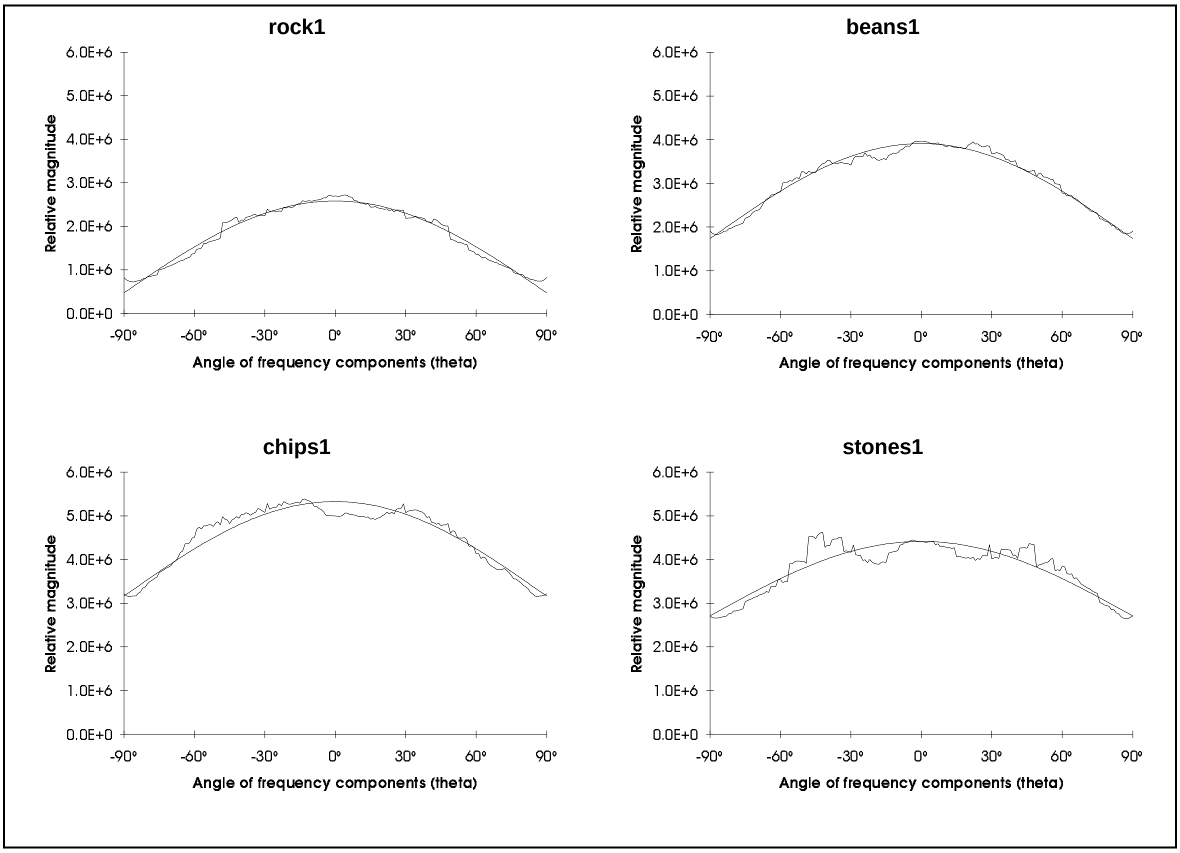

Figure 3.26 - Polar plots, and best fit cosines, of the textures beans1, chips1, and stones1 (ι = 0°).

Table 3.2 and figure 3.26 show that *stones1* does indeed exhibit a higher platform than *rock1*, as do *chips1* and *beans1* — supporting the suggestion that the platform height is related to surface variance. This support however is tentative given the small sample and lack of quantitative surface height data. What is clear however, is that all four image textures exhibit distinct directional characteristics which "follow" the angle of tilt. These empirical results therefore

(i) show that image texture directionality is not an intrinsic characteristic, and

(ii) support the $\cos(\theta - \tau)$ relationship between illuminant tilt and image texture, but show that it should be more accurately modelled by adding an additive term to account for the raised cosine effect. That is it shows that the tilt component should be modified to :

$$F_\tau(\omega,\theta) = m_\tau \cos(\theta - \tau) + b_\tau \tag{3.8}$$

d) Radial shape - an intrinsic characteristic ?

The above shows that the directional characteristics of image texture are not independent of illuminant tilt — as predicted by the image model presented in chapter 2. This model also predicts that radial shape of magnitude spectra is an *intrinsic* property (see figure 2.5). If this is indeed the case it will be independent of illuminant tilt. Thus radial plots of image texture will show that variation in tilt changes the level but not the form of the log-log radial response. Figure 3.27(a) shows the response of *rock1* image texture to variation in illuminant tilt (τ). It contains radial sections through the periodograms at an angle θ = 0°. These sections show that, as predicted by the image model (2.14), the magnitudes reduce as τ deviates from θ, and that the plots are similar to each other in shape although they converge towards each other at the Nyquist frequency.

Figures 3.27 (b), (c) and (d) show radial sections of the other three isotropic textures. These plots together with those at other values of θ all show similar results — gross radial shape is maintained but the plots converge as the Nyquist frequency is approached. That is the gradients of sections (particularly of *beans1* and *chips1*) are dependent upon the tilt angle of the illuminant. Hence, contrary to predictions derived from the image model (2.14), estimates of the power roll-of factors of these textures would not in this instance provide a feature which is independent of illuminant tilt angle.

a) rock1 *b) beans1* *c) chips1* *d) stones1*

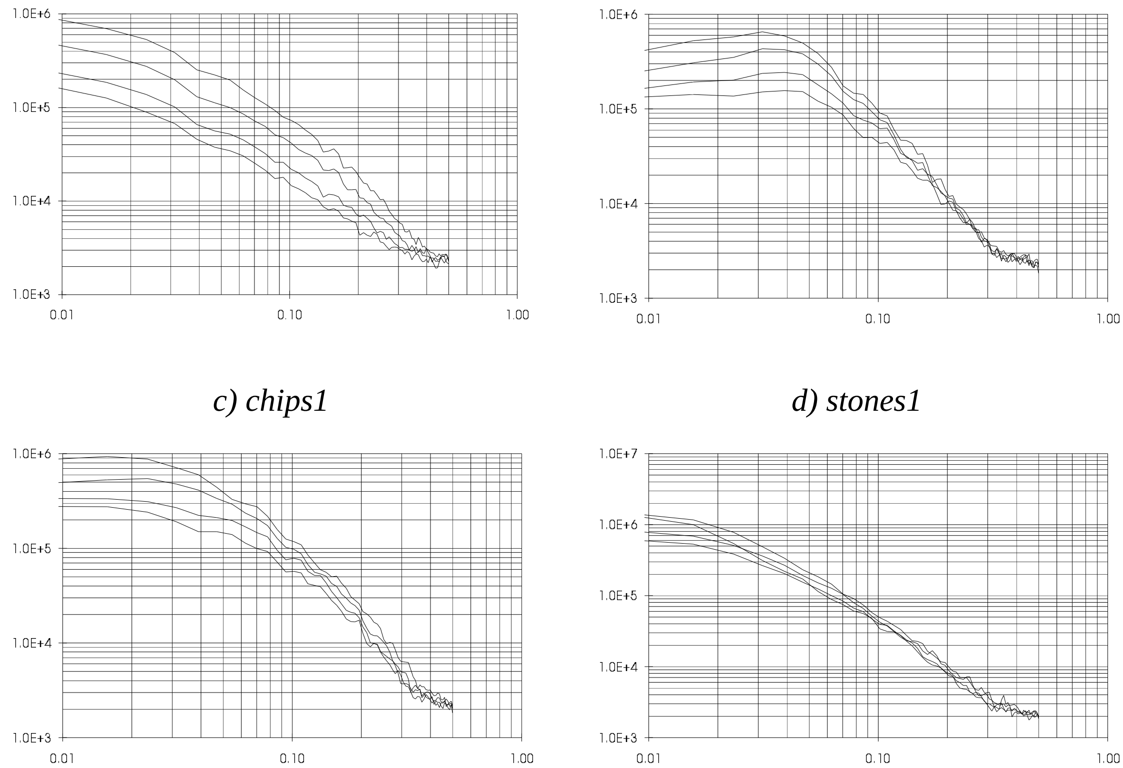

τ = 0° (top traces), τ = 30° (2nd top traces), τ = 60°(2nd bottom traces), τ = 90° (bottom traces)
θ = 0°, σ = 50°.
Axes as for figure 3.24

Figure 3.27 - Effect of tilt on radial shape of magnitude spectra (axes as previous figure)

3.2.5. Summary of tilt response investigation

This section has presented the results of an investigation into the effects of variation in illuminant tilt angle, on image texture; through the use of simulation and laboratory experiment. To summarise :

- Results from simulation and experiment show that the directional characteristics of image texture are not intrinsic — but that they are dependent upon illuminant tilt.
- The linear image model (2.14) predicts a pure cosine relationship : $\cos(\theta - \tau)$ but the results of simulation and laboratory experiment show that a raised cosine : $F_\tau = m_\tau \cos(\theta - \tau) + b_\tau$ — is more appropriate for textures with larger slope angles.

- The shapes of radial sections of the test textures at different illuminant tilt angles are similar, although convergence of the sections towards the Nyquist frequency was observed. Thus in this instance estimates of the power roll-off factor β are not independent of the tilt angle of the illuminant.

3.3. The slant angle response of image texture

The two preceding sections of this chapter have investigated the response of image texture to variations in surface relief and illuminant tilt angle (τ). The results presented support the first two parts (F_s and F_τ) of the image model of topological texture (2.14). The third component of the model concerns the response of image texture to changes in illuminant slant angle (2.17) :

$$F_\sigma = \sin \sigma$$

This implies that, as the angle the illuminant vector makes with the vertical is increased, the whole magnitude spectrum is uniformly amplified by a factor equal to the sine of that angle.

The aims of this section are therefore :

(i) to assess the validity of the slant response predicted by (2.17), and

(ii) to further investigate the intrinsic nature of the radial shape of image texture magnitude spectra.

As in the previous section, simulation was used to gain an insight into the effect of high slope angles and shadowing, while laboratory experiment provided results with real textures. The intrinsic nature of PSD radial shape is discussed again here — from the perspective of slant angle response; and the section finishes with a summary.

3.3.1. Low slope angles

Synthetic images of texture were generated as described in section 3.2.1 using low average slope angles (height scaling factor $S = 1$). For these simulations the illuminant's tilt (τ) was kept constant at 0°, while the slant angle (σ) was varied in 10° steps between

10° and 80°. Figures 3.28 and 3.29 show samples of four of the resulting images and their magnitude spectra.

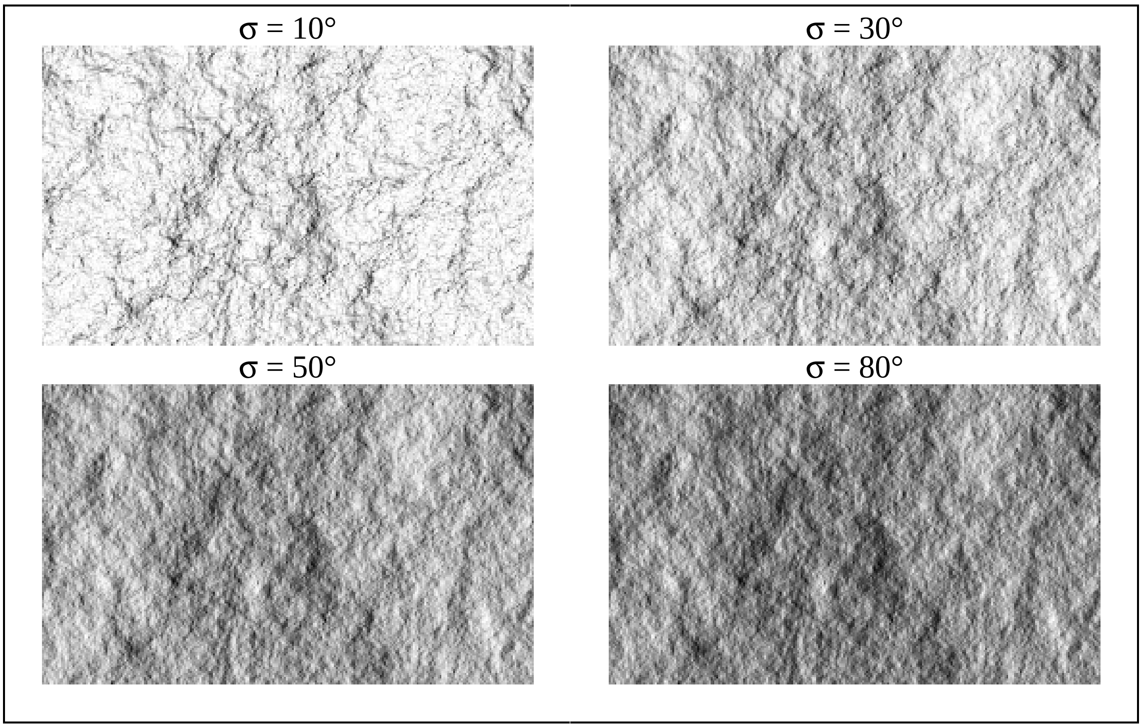

Figure 3.28 - Samples of intensity images — showing the effect of changing σ .

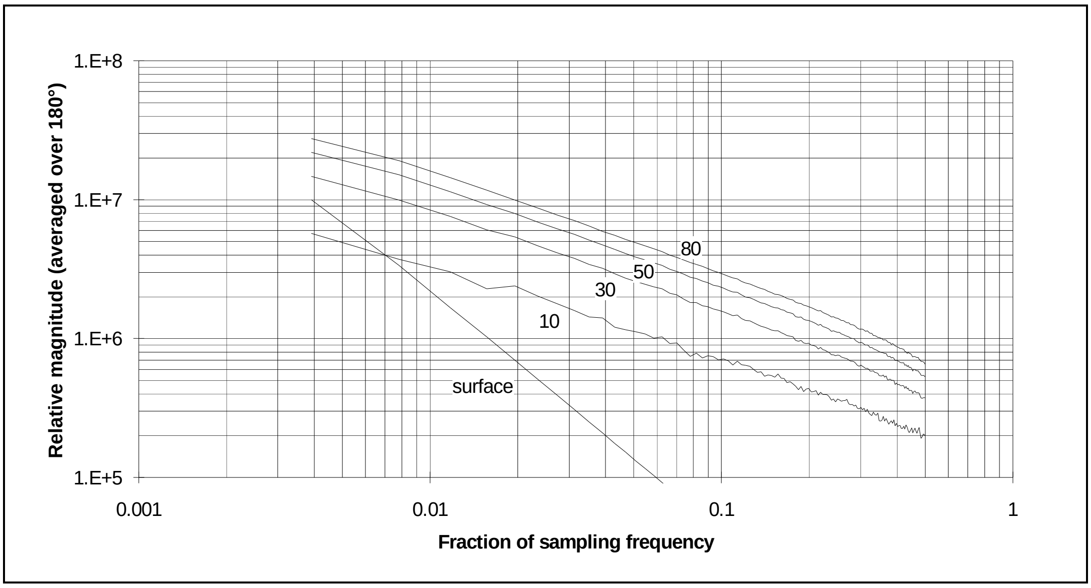

Figure 3.29- Frequency spectra of surface height map and intensity images for σ= 10^{c}, 30^{c}, 50^{c}, and 80^{c}.

As predicted by the image model (2.14) the slope of the above spectra do not change (i.e. the power roll-off factor β_I remains constant) but the variance of the image texture does increase with slant angle (σ). In order to establish whether equation (2.17) represents the slant angle response, magnitudes of frequency components (at $\omega = 0.12\omega_S$) were estimated for each image[4]. These estimates were calculated using the least squares fit of a straight line to the log-log magnitude spectra. They are plotted against sin σ in the graph below.

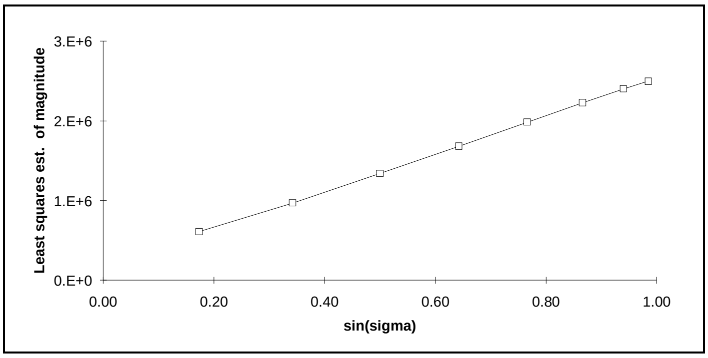

Figure 3.30 - Slant angle (σ) response showing the sinσ relationship (for σ=10°,20° ...80°).

The above graph shows that the simulated magnitude response is a linear function of sin σ as predicted by the image model (2.14).

3.3.2. Large slope angles and shadowing

The simulations were repeated for a variety of surface variances in order to test the applicability of the slant angle relationship for larger slope angles. Figure 3.31 shows that the relationship remains linear, although it is obvious that the y-intercept constant increases with surface variance.

Unfortunately this linear relationship does not continue to hold once shadowing has been introduced. Figure 3.32 shows images at the same slant angles as before, but for a surface variance of 112 (S = 4) and with the addition of shadowing.

[4]Where ω_S is the sampling frequency.

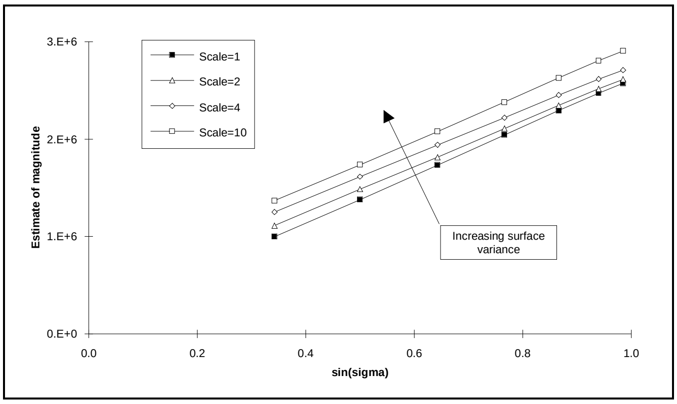

Figure 3.31 - Effect of power (scale = 1, 2, 4 & 10) on slant angle response

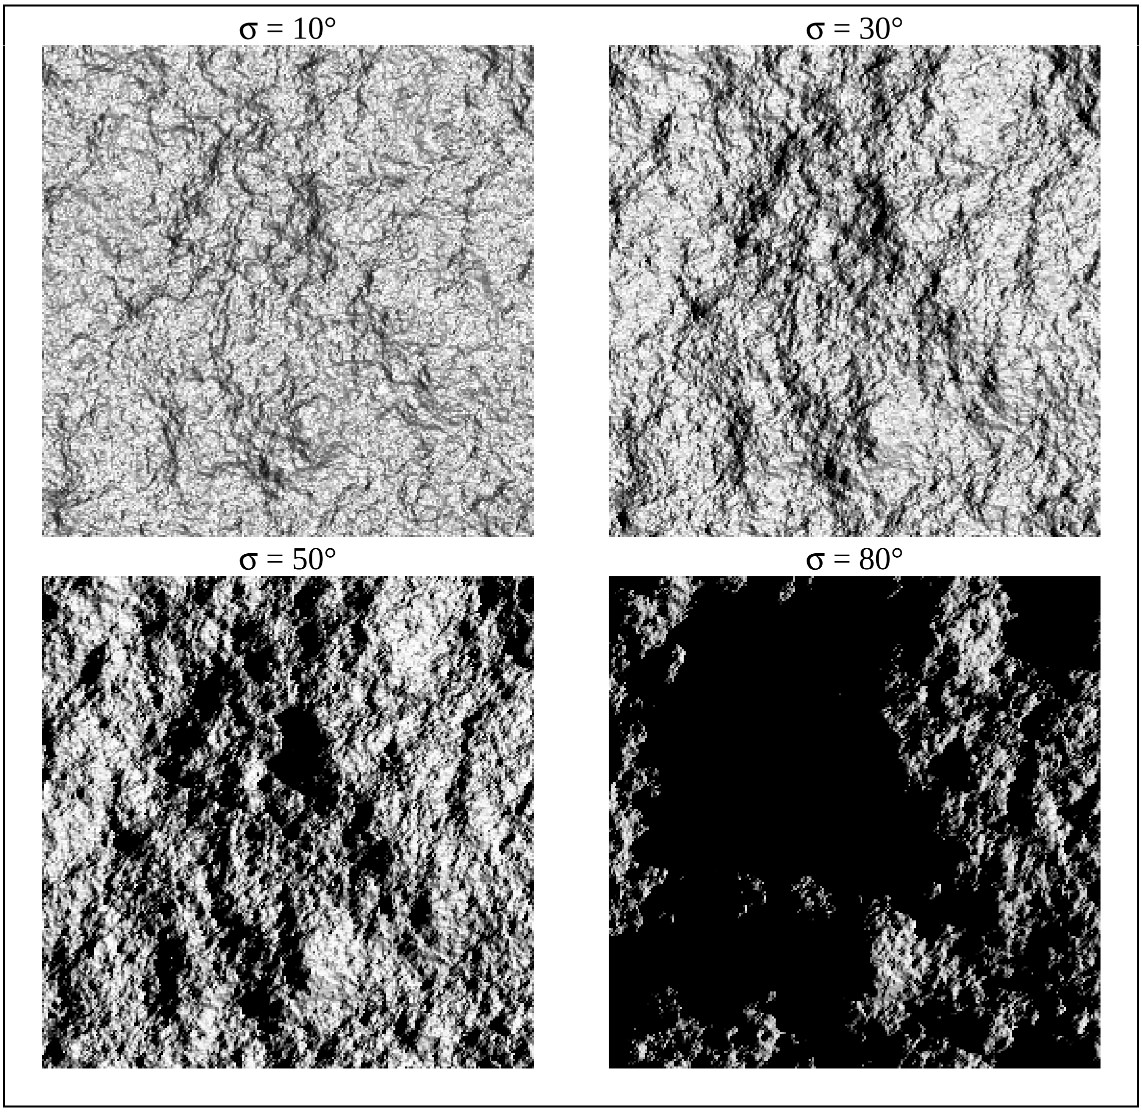

Figure 3.32 - Intensity images showing variation due to change of slant angle for a height scaling factor S = 4

As the graph below shows adding shadowing to the simulation severely distorts the linear relationship.

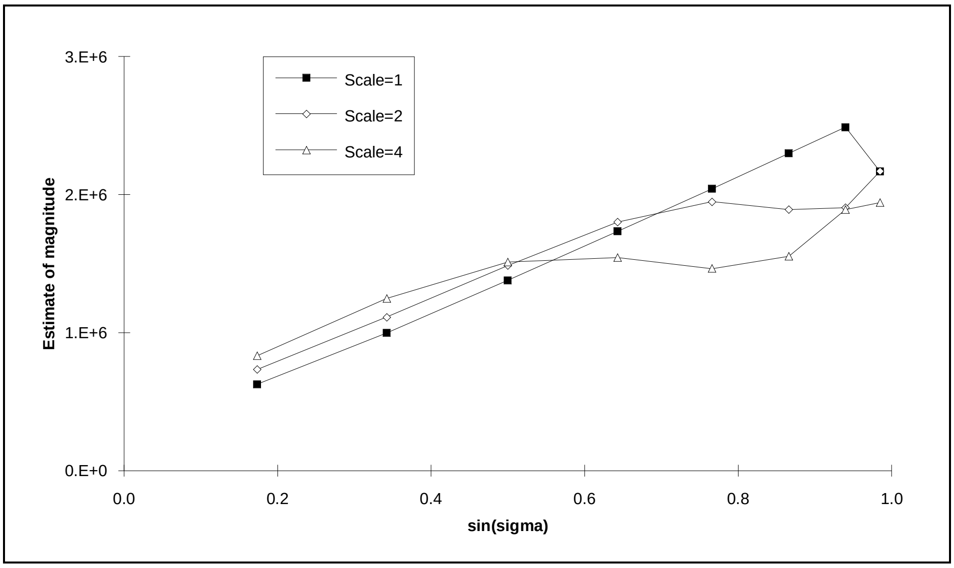

Figure 3.33- Effect of shadowing at various powers (scale = 1, 2 & 4) on the slant angle response.

For each of the height scaling factors shown, a significant reduction in magnitude occurs when the area in shadow exceeds 1-2%. Thus the slant angle response is severely modified by even slight shadowing.

In conclusion therefore, simulation results predict that "sin σ" slant response holds while the degree of shadowing is small, but that it is severely affected by even small amounts of shadowing.

The following section therefore investigates this relationship using four real textures.

3.3.3. Experimental results : slant response

The four textures used in the tilt angle experiments; *rock1*, *beans1*, *chips1*, and *stones1*, were imaged as before, except that illuminant tilt was kept constant at $\tau = 0°$, and slant was varied in 10° steps between 10° and 80°. Four samples of the resulting intensity images and their magnitude spectra are shown in figures 3.34 and 3.35.

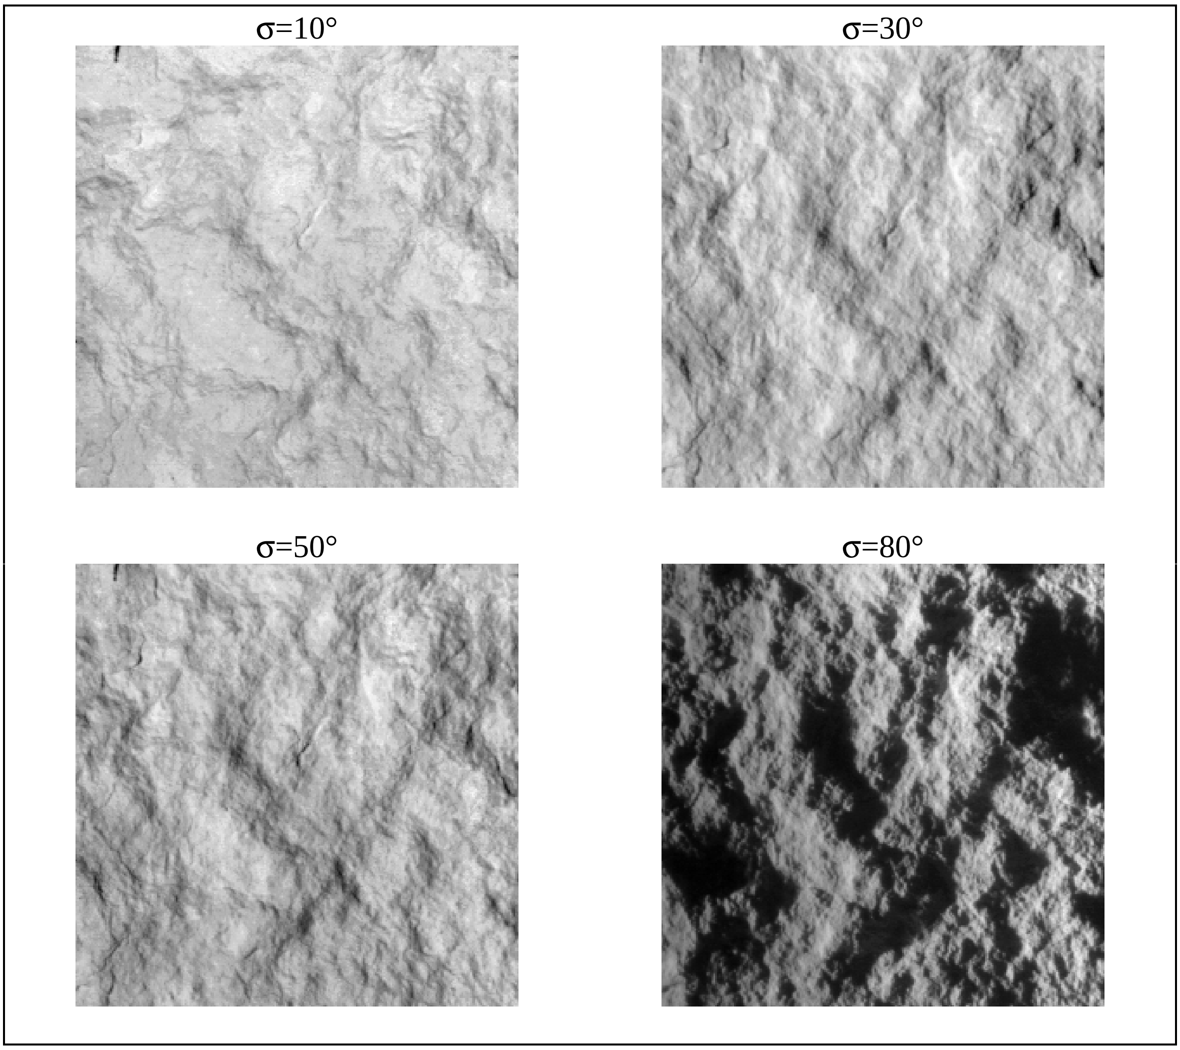

Figure 3.34 - Intensity images of rock1 showing variation with slant angle

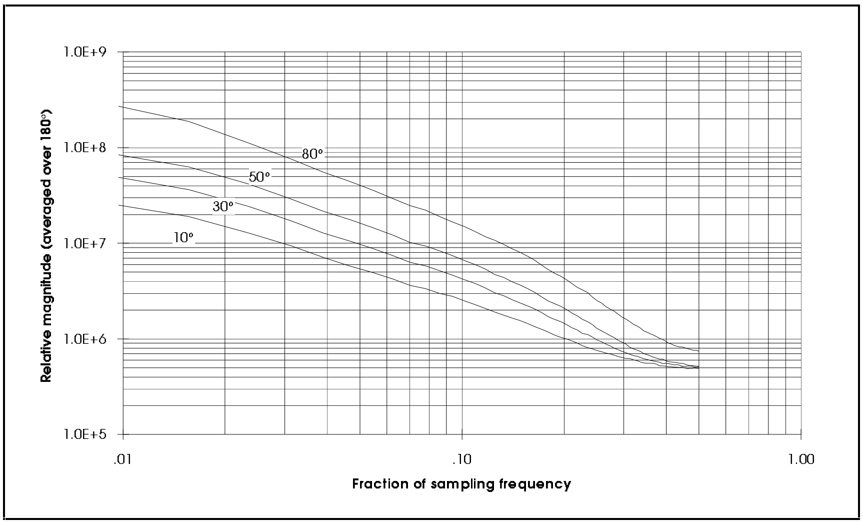

Figure 3.35 - rock1 : slant angle response ($\sigma = 10^{c}$, 30^{c}, 50° & 80^{c})

Figure 3.35 of average radial sections of *rock1's* magnitude spectrum shows that image variance does increase with slant angle. It is however difficult to assess whether or not it follows the sin σ relationship (2.17) derived in chapter 2. Estimation of the magnitude at a particular frequency via a straight line approximation is not appropriate here, as radial sections of spectra of the test textures are not straight lines. Hence a simple alternative was employed : the average of the coefficients in the range $\omega = 0.05\omega_S$ to $0.2\omega_S$ was taken. These magnitude estimates were plotted as before against sin σ, and are shown below in figure 3.36.

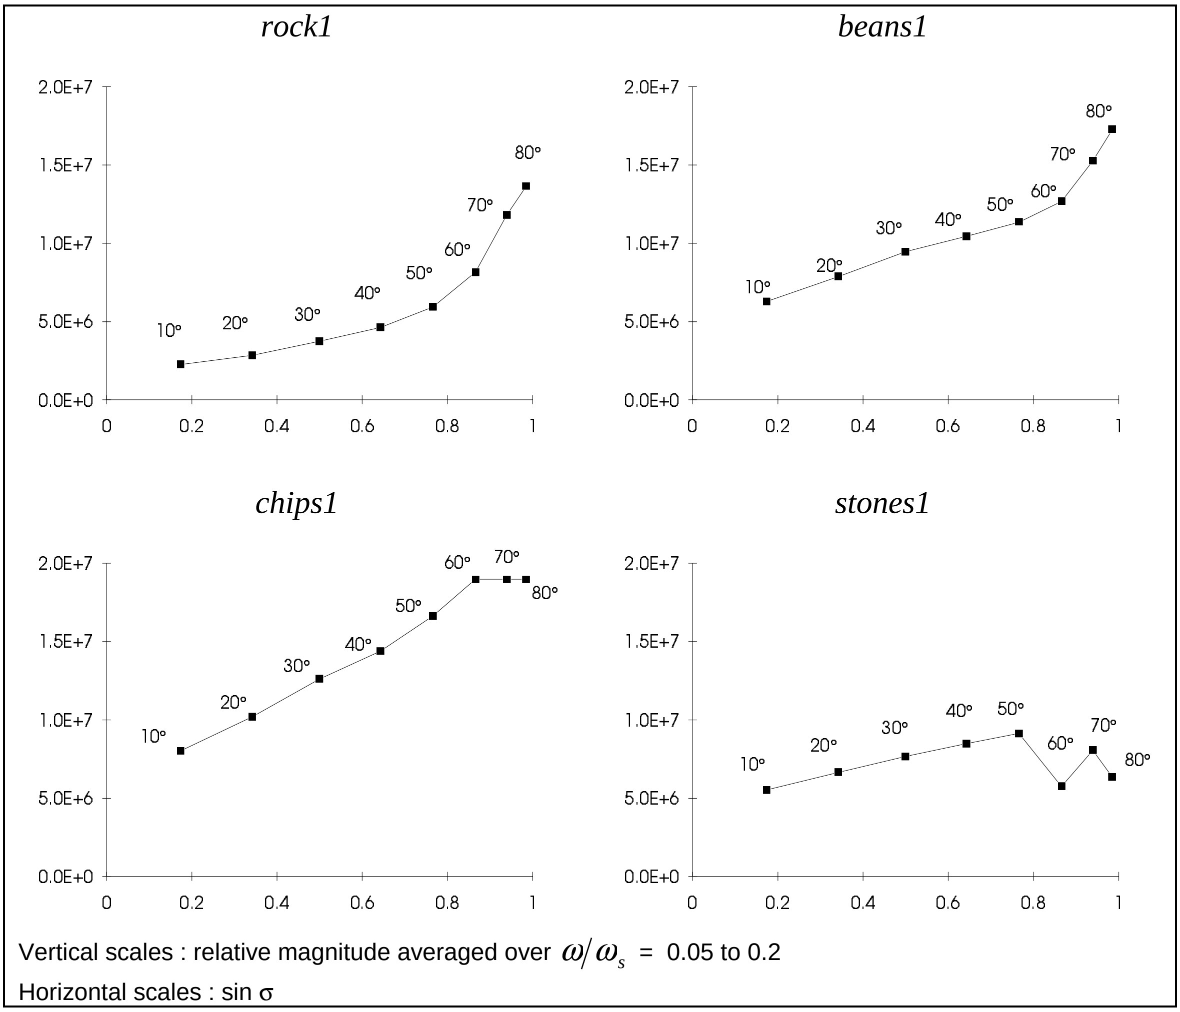

Figure 3.36 - Magnitude estimates vs. sin (σ)

Clearly the graphs above do not display a linear relationship over the entire range of slant angle. However, for the lower values of slant ($\sigma \leq 50°$), where the shadowing has less effect, the graphs do show a sin σ relationship.

3.3.4. Radial shape - slant angle response

The above has shown that, with respect to the test textures, the values of magnitude spectra are *dependent* upon slant angle, and that a sine relationship holds for slant angles of 50° or less. The image model developed in chapter 2 predicted that radial shape is an *intrinsic* property of texture, and therefore *independent* of illuminant slant σ. Figure 3.37 allows the intrinsic nature of this characteristic to be assessed for the four test textures. It shows radial sections through the two dimensional magnitude spectra, at θ = 0° for four values of illuminant slant. It can be seen that the shapes of the graphs do not change significantly with variation in illuminant slant. However, as with the tilt angle response, the plots again converge towards the Nyquist frequency, and except for *stones1,* their gradients are not independent of σ. Hence estimates of the power roll-off factor β_I, are unlikely to provide a texture measure which is purely a function of the surface relief.

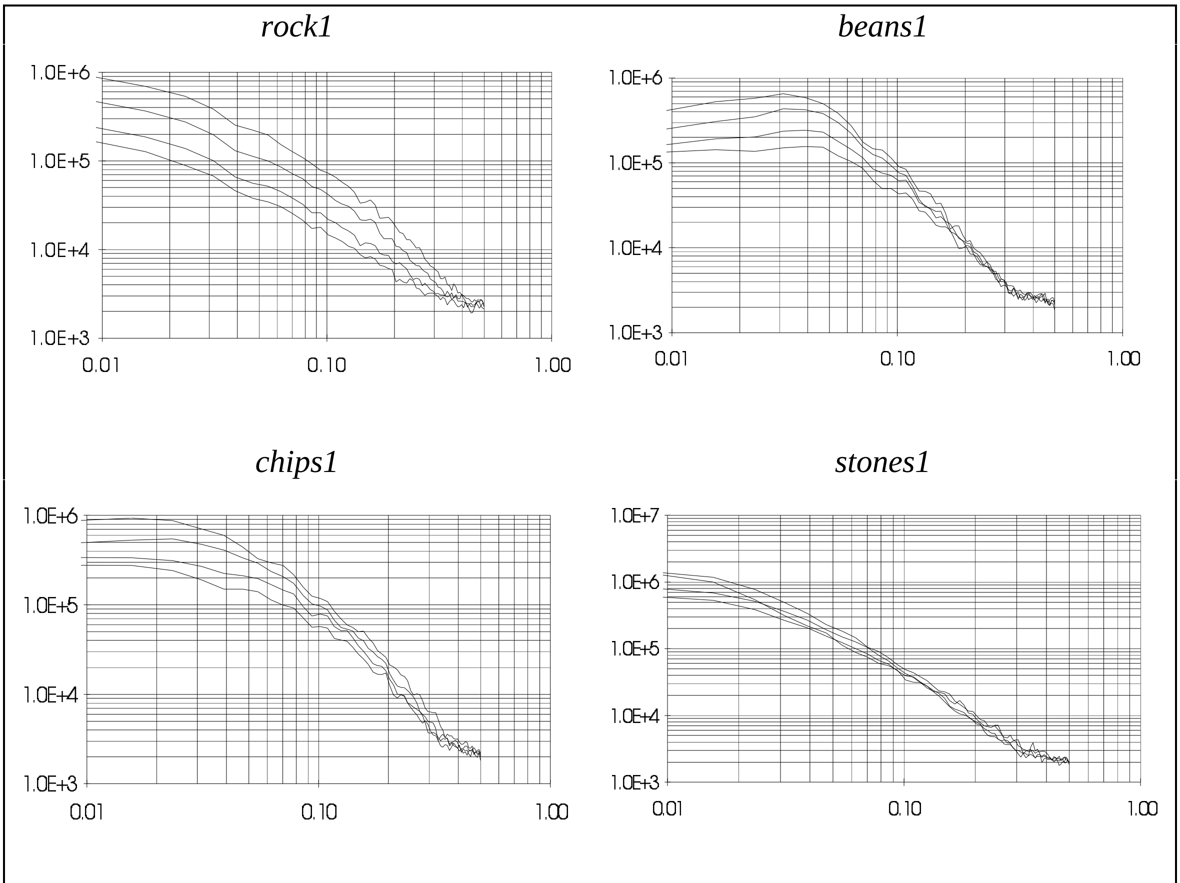

σ = 10°(bottom trace), 20° (2nd bottom), 40° (2nd top) & 60°(top trace),
θ = 0°, vertical scales : relative magnitude, horizontal scales : fraction of sampling frequency.

Figure 3.37 - Radial shape : slant angle (σ) response.

3.3.5. Summary of slant response investigation

This section has investigated the slant angle (σ) response of image texture.

- The results of simulation and laboratory experiment have shown that image magnitude spectra are not independent of the illuminant's slant angle.
- In simulation, shadowing severely affected the predicted "sin σ" relationship.
- Laboratory experiments have shown that the slant angle responses of the four test textures, approximates a linear function of sin σ for slant angles of up to 50°.
- Laboratory experiments have also shown that the gross radial shape of magnitude spectra of the four test textures, is unaffected by illuminant slant. However, the gradients of these spectra (and hence the power roll-off factors) are not independent of illuminant slant.

3.4. Conclusions

This section summarises the investigations reported in this chapter and briefly assesses their likely impact on texture classification and segmentation.

Chapter 2 presented an image model of topological texture due to Kube and Pentland [Kube88]. This model is important to texture classification and segmentation as it predicts that many texture features will be affected by changes in illuminant direction. However, the model was derived assuming that slope angles are low, and shadowing was ignored. Thus the purpose of this chapter was to investigate the model's validity particularly with regard to these two aspects.

Chapter 2 divided the model into three parts corresponding to

(i) the response of image texture to changes in surface relief,

(ii) the response of image texture to changes in the tilt angle of the illuminant (τ), and

(iii) the response of image texture to changes in the slant angle of the illuminant (σ).

Hence this chapter reported results from three investigations; one for each type of response. The main conclusions from each of these are repeated below :

a) The response of image texture to changes in surface relief

The investigation into the effect of surface relief on image texture used an isotropic fractal model of topological texture. For such surfaces the radial shape of the surfaces' PSD plot is of the form $f^{-\beta_H}$ and the image model (2.14) predicts that the radial shape of the image's PSD plot will be of the form $f^{-\beta_I}$, where $\beta_I = \beta_H - 2$ (the β relationship). Results showed that in simulation :

- the β relationship is representative over a range of surface variances, and
- the β relationship is still valid when shadowing occurs.

These results therefore, also support the surface response component of the image model (2.14), from which the β relationship was derived.

b) The response of image texture to changes in the tilt angle of the illuminant

The second component of the image model (2.14) predicts that the tilt response of a texture is of the form :

$$F_\tau = \cos(\theta - \tau)$$

Simulation and laboratory experiment were used to investigate this response :

- Results from simulation and experiment show that the directional characteristics of image texture are not intrinsic, but are dependent upon illuminant tilt.
- The results of simulation and laboratory experiment, show that a *raised* cosine (3.8) : $F_\tau = m_\tau \cos(\theta - \tau) + b_\tau$, rather than the straight cosine relationship above, provides a more accurate representation of the tilt response.

c) The response of image texture to changes in the slant angle of the illuminant

The third component of the image model (2.14) predicts that the slant response of a texture is of the form :

$$F_\sigma = \sin \sigma$$

As for the previous response both simulation and laboratory experiment were used in the investigation. They showed that :

- the variance of image texture is not an intrinsic characteristic as it is dependent on the slant angle of the illuminant,
- that shadowing severely affects the predicted sin σ relationship, and

- that for the four test textures the slant angle response follows a sine law for values of slant angle less than or equal to 50°.

d) The intrinsic nature of PSD radial shape

Chapter two's linear image model (2.14) implies that the radial shape of power and magnitude spectra are independent of the illuminant's direction. This characteristic was investigated during slant and tilt experiments. Both showed that the gross shapes of radial sections of the test textures are invariant to the direction of the illumination. However their gradient is affected by variation in the illuminant vector. Hence estimates of the power roll-off factor would not provide a texture feature which is invariant to changes in the orientation of the illuminant.

3.4.1. Implications for texture classification

Many of the feature sets surveyed in chapter 4 contain directional texture measures. In addition some are clearly a function of image variance (see chapter 4). Hence two of the most important of the above conclusions are that

(i) the directionality of image texture is not solely a function of surface directionality, but that it is also a function of illuminant tilt, and

(ii) that variation in illuminant slant, can also affect image variance.

Thus classification accuracy may well be reduced if the direction of the illumination either (a) changes between training and classification sessions, or (b) varies over a scene due, for instance, to the proximity of the lighting source.

The purpose of the next chapter therefore, is to review and choose sets of texture features for investigation as to the effects of variation of illuminant direction.

Chapter 4

Texture features : review and selection

The two preceding chapters have used theory, simulation, and laboratory experiment, to investigate the way in which changes in illuminant direction affect image texture. For the test sets employed, it has been shown that variations in either illuminant slant or tilt affect image texture. It was suggested that normalisation may compensate for changes in the former but not the latter; as variation in tilt affects the directionality of image texture. Since directional features are used in texture classification and segmentation schemes; it is to be expected that variation in tilt may affect the performance of some of these schemes.

In reviewing texture features for use in classification and segmentation schemes this chapter therefore has two main objectives :

(i) to identify research on the effect of variation in illuminant direction on texture classification and segmentation, and

(ii) to select three sets of feature measures for further investigation as regards changes in illuminant slant and tilt.

However, it is not practical to provide an exhaustive survey of all texture measures here. Given the extent of the literature such a task is beyond the scope of this thesis. Rather this chapter reviews some of the more popular techniques. Concerning point (ii) above, the criteria used for the selection of the features were :

(a) popularity in the literature,

(b) ease of implementation and use, and

(c) reported performance.

Before describing feature measures in earnest, the meanings of three terms necessary to the following discussion : *segmentation, classification* and *feature measure,* are defined. The review itself starts by considering survey papers, which divide the subject into statistical and structural approaches. Only the former will be reviewed here. Within statistical methods two further groupings can be discerned. They are (a) model-based approaches (i.e. based upon parameter estimation techniques) and (b) non-model-based methods. Descriptions of these two feature measure types are followed by a resume of "rotation invariant" feature measures : the motivation for this discussion being that the rotation of subject texture under fixed lighting, would normally imply an effective change in illuminant vector relative to the texture. Finally the conclusion summarises the literature reviewed, with particular emphasis on the consideration given to lighting variations, and it identifies three types of feature measure for further investigation.

4.1. Definition of segmentation, classification and feature measure.

Before discussing the various texture features it is helpful to clarify what the terms *segmentation*, *classification* and *feature measure*, as used in this thesis, refer to. Texture *segmentation* is used to refer to the process of dividing an image up into homogeneous regions according to some homogeneity criteria. It is therefore intimately concerned with establishing the boundaries between these regions without regard to the type or class of the regions. For brevity the term *segmentation* on its own will normally be used to describe this process.

Texture *classification* refers to the process of grouping test samples of texture into classes, where each resulting class contains similar samples according to some similarity criterion. If the classes have not been defined a priori, the task is referred to as *unsupervised* classification. Alternatively, if the classes have already been defined (normally through the use of training sets of sample textures) then the process is referred to as *supervised* classification. In the following text nearly all the texture classification tests reported are of the supervised type and this process will be referred to simply as *classification.* Unsupervised classification will therefore be specifically identified as such. Note that classification tests may be performed on separate samples of texture, in which case the samples are presented as a number of separate, normally square, images and segmentation is not required. Alternatively, a single image containing multiple textures may be presented, requiring segmentation prior to classification. Note that if classification is performed on a pixel by pixel basis within a single image then, as a by-product, segmentation also occurs.

Before either segmentation or classification can take place, some homogeneity or similarity criterion must be defined. These criteria are normally specified in terms of a set of *feature measures,* which each provide a quantitative measure of a certain texture characteristic. These *feature measures* are alternatively referred to here as *texture measures* or just simply *features.* Groups of feature measures assembled for segmentation or classification purposes are often referred to as *feature vectors.*

Note that when the performance of feature measures are compared, it is misleading to compare classification and segmentation accuracies. The former normally refers to the percentage of correctly classified texture samples or regions, while the latter may refer to the number of correctly identified *pixels.*

4.2. Surveys

Haralick provided the classic survey of texture measures [Haralick79]. He listed and described a number of texture extraction methods which he divided into two types : structural and statistical, as did Wechsler [Wechsler80]. More recently Van Gool et al produced an excellent survey of texture analysis [VanGool85]. They again divided the field into structural and statistical camps. The former use primitives to describe texture elements and placement rules to describe the spatial relationship between elements. This approach is better suited to textures that exhibit a regular macro-structure, and will not be discussed further. The statistical approaches are better suited to micro textures, and Haralick identified techniques based upon auto correlation functions, frequency domain analysis, edge operators, grey-level co-occurrence matrices, grey-level run lengths, and autoregressive models. The taxonomy of statistical techniques due to Van Gool et al, is similar to Haralick's, but it also includes the use of filter masks such as Laws' energy features, and grey-level sum and difference histograms. In addition it provides a summary of reported performance.

The surveys referred to above were performed in the late seventies and early eighties since which there has been an explosion of interest in model-based techniques (Markov fields, fractals etc.). These are well reviewed in a recent survey by Reed and du Buf [Reed93], which also covers feature-based methods (including statistical approaches) and structural methods.

The following review is therefore divided into two main groupings : model-based and non-model-based features.

4.3. Model-based features

A number of random field models (i.e. models of two-dimensional random processes) have been used for modelling and synthesis of texture. If a model is shown to be capable of representing and synthesising a range of textures, then its parameters may provide a suitable feature set for classification and/or segmentation of the textures. For a model based approach to be successful, there must exist a reasonably efficient and appropriate parameter estimation scheme, and the model itself should be parsimonious, i.e. use the minimum number of parameters. Popular random field models used for texture analysis include fractals, autoregressive models, fractional differencing models, and Markov random fields. These will now briefly be reviewed. A more extensive review of these approaches may be found in [Ahuja81] and [Reed93].

4.3.1. Fractal models

Fractals [Mandelbrot83] have, as discussed earlier, been used very successfully to synthesise natural looking textures [Voss88] [Saupe88]. Their use for synthesis together with their ability to characterise "roughness" [Pentland84] make their major parameter, fractal dimension, a natural candidate as a feature measure of texture. Many researchers have estimated the fractal dimension and used this directly as a texture measure. Such estimates are obtained either in the frequency domain, by estimating the gradient of the log-log plot of the power spectrum, or from the spatial domain by a variety of methods

[Voss88]. Note however, that as fractal dimension describes *scaling* behaviour, it is necessary to perform measurements over at least two scales, and that one would expect that the wider the variation in scales the more accurate would be the estimation procedure. Accurate estimation of fractal dimension therefore seems to be at odds with the accurate determination of texture boundaries — however, this is a trade-off that all texture segmentation schemes must make.

Pentland reported one of the earliest uses of fractal dimension estimates for segmentation purposes [Pentland84]. He used the power spectrum method to provide an omnidirectional estimate of the fractal dimension of 8x8 pixel blocks, which he used to segment a variety of indoor and outdoor scenes. From the results presented it appears that the scenes have all been coarsely segmented into textured and non-textured regions, something which could not have been achieved with straight grey-level thresholding.

Medioni and Yasumoto [Medioni84] also used a single omnidirectional fractal dimension estimate as a feature measure. They tried to segment an image containing multiple textures and obtained unsatisfactory results. They commented that *"it* (fractal dimension) *suffers the drawbacks associated with any single feature measurement space: it describes one aspect of texture and therefore can only separate textures which differ enough in roughness".* Keller and Chen [Keller89] similarly state that *"fractal dimension alone is not sufficient to characterise natural textures".*

Two techniques have been used to enhance the classification power of fractal dimension based methods. Firstly, additional parameters such as "lacunarity" have been utilised. Secondly, directionality, a key feature of most texture analysis schemes, has been employed by relaxing the assumption of isotropy and providing directional estimates of fractal dimension.

The term lacunarity was derived by Mandelbrot from the Latin term for gap (lacuna) and he used it to describe the size of holes in images of galaxies [Mandelbrot83]. Linnett [Linnett91a] likened this characteristic to "structure" in texture, which is dependent on phase information [Clarke92]. Keller and Chen [Keller89] used a measure of lacunarity together with an omnidirectional estimate of fractal dimension for segmentation purposes. They found that the use of this additional parameter, considerably improved the results achieved with a test image, consisting of a montage of Brodatz textures [Brodatz66].

Directionality is commonly exploited for texture classification and segmentation. It is not surprising therefore, that when Pentland [Pentland84] extended his study to compare the performance of his method against that of Laws, he used estimates of fractal dimension in x and y directions. Pentland reported a classification accuracy of 84.4% on a Laws test image (a Brodatz texture montage) which compared well with other techniques. Mosquera et al [Mosquera92] estimated fractal dimension in four directions (vertical, horizontal and the two diagonals) using a spatial domain method based on a 12x12 kernel. They achieved good segmentation results on synthetic and Brodatz textures.

Peleg [Peleg84] extended the "sausage" method of measuring the length of a fractal curve [Mandelbrot83] to produce a "blanket" method of measuring the area of a fractal surface. In both cases the way in which the measurement (length L or area A) scales with the "measurement yardstick ε" is an exponential function of the fractal dimension D. Thus D can be estimated from the gradient of the log-log plot of A against ε. Peleg used an iterative method for the calculation of the area A at different values of ε — he simply calculated the position of the next blanket's surfaces by adding a radius onto the previous blanket's surfaces. In this manner he was able to generate 50 blankets, from which he obtained 48 estimates of the log-log gradient (each from a set of three blankets). The 48 gradients were then used as features to successfully classify a set of Brodatz textures.

Peleg also described the possible use of directional versions of his "blanket" based feature, a technique which Linnett and his colleagues [Linnett91a] [Linnett91b] [Linnett93] used to great effect in his segmentation scheme. As well as using directional operators Linnett

(i) used the blanket thickness as a feature measure directly, rather than the estimates of fractal dimension,

(ii) used the thickness at each position in the image for spatially accurate segmentations,

(iii) used only a low number of blanket iterations (typically two or three) thereby reducing the computational load considerably,

(iv) used a moving window averaging filter to improve the robustness of the classification (by reducing feature variances), and

(v) used an iterative statistical classification method [Linnett91a].

Linnett achieved impressive results on side scan sonar images of the sea bed and a Brodatz texture montage (4.3% classification error).

Clarke [Clarke92] further extended Linnett's techniques to provide an elegant method of *rotation invariant* classification. His scheme transforms the feature space of each segmented region using principal components analysis (PCA). The analysis is performed separately on each region in turn and the resulting principal components are used as features for classification.

The quality of these results inspired the author to apply Linnett's scheme to underwater images [Chantler91] and to embark on this research.

4.3.2. Autoregressive models

Autoregressive models have been used for spectral estimation [Marple87], coding [Kashyap80], segmentation and classification [Khotanzad87], and image restoration [Chellappa82].

A *time series* autoregressive model is a random process model in which the current value of the output is expressed as the sum of its mean value, the current value of a white noise process, and a linear aggregate of previous output values [Box76]. The number of output values used is known as the *order* of the model (p). An autoregressive model therefore has $p + 2$ parameters: p coefficients, the mean, and the variance of the white noise. These parameters may be estimated using either least square error or maximum likelihood techniques [Khotanzad89]. Autoregressive models have been used to model images as random fields (two-dimensional random processes) by a number of researchers [Kashyap83] [Kartikeyan91] [Mao92]. In the two-dimensional spatial case the "previous values" of the time series process are replaced by the grey values of local neighbourhood pixels. Unlike the temporal case there is normally no *preferred direction* in a lattice and neighbourhoods are therefore normally defined to consist of variables both "before" and "after" the variable being modelled, i.e. they are *non-causal* (two-sided). Again the parameters may be estimated either by using least square error or maximum likelihood techniques.

Khotanzad and Kashyap [Khotanzad87] used "simultaneous autoregressive models" for texture classification. They selected the order of the models to use via texture synthesis. They fitted simple local neighbourhood models to the test textures and the resulting parameter sets were used to synthesise textures. If the synthesised textures were dissimilar to one another, then their parameters were deemed to have good discriminatory powers over the test textures. However, if the converse were true, the process was repeated with a higher order model.

Kartikeyan & Sarkar [Kartikeyan91] reported a classification scheme in which they first identified the most suitable model parameters for *each* training class and then estimated the value of the parameters themselves. Thus each training class has its own feature space, consisting of the parameters of its model. Classification was performed by calculating a set of feature vectors (one for each of the feature spaces) for a test texture and determining likelihood of the texture belonging to a class by using the corresponding class feature space. This method achieved a miss-classification rate of "about 2.4%" on a set of four Brodatz textures.

As far as parameter estimation is concerned, Khotanzad and Chen [Khotanzad89] found little difference between least squares and maximum likelihood methods. They used a six parameter autoregressive model and edge detection in feature (parameter) space for segmentation of natural textures.

Kashyap and Khotanzad [Kashyap86] developed a *rotation invariant* classification scheme which uses two autoregressive models. Firstly they use a "circular symmetric autoregressive model" in which all the weights for the local neighbourhood are lumped together to give an isotropic (directionally insensitive) parameter. Secondly, a conventional autoregressive model is used to provide a measure of "directionality" : this feature is essentially the maximum difference between right-angled pairs of neighbourhood parameters e.g. (1,1) and (1,-1). Kashyap and Khotanzad tested their scheme on twelve Brodatz textures at seven different rotations and obtained an average classification accuracy of 91%.

Mao & Jain [Mao92] developed a similar *rotation invariant* autoregressive model. It also uses the sum of unit circle based neighbourhood grey-levels, but in addition extends the order of the model to take in parameters based on wider radius circles. A directional measure is not used. Instead, to improve classification accuracy, a multi-resolution approach was adopted, in which the parameter estimation is repeated at different scales. By using a second order model at four different resolutions 100% classification accuracy was achieved. It should be noted however that, as in other "rotation invariant" tests, the test set was created by rotating *images* rather than rotating the natural textures themselves.

4.3.3. Fractional differencing models

Kashyap and his colleagues have extensively investigated the use of *fractional differencing models* (also termed *long correlation models*) for modelling, synthesis, classification, and segmentation, of texture [Kashyap84] [Kashyap89] [Choe91a] [Choe91b]. The one-dimensional fractional differencing model was suggested by Hosking [Hosking81] as a generalisation of Box and Jenkins ARMA(p,d,q) model, where p,d,q are the orders of the autoregressive (AR), differencing, and moving average (MA) parts of the model respectively [Box76]. The generalisation is a relaxation on the differencing parameter d (which is normally a low valued positive integer) to allow it to be real valued (i.e. fractional). Hosking defines the *fractional difference operator* ∇^d of order d using the binomial series :

$$\nabla^d = 1 - dB - \tfrac{1}{2} d(1-d)B^2 - \tfrac{1}{6} d(1-d)(2-d)B^3 \ldots\ldots$$

where B is the backward shift operator.

Thus for positive fractional values of *d* the ARMA(0,*d*,0) process is in fact equivalent to an infinite order AR(autoregressive) model. This explains its "long memory" characteristics [Choe91a].

Kashyap and Eom used this model to develop a texture segmentation scheme. The model's parameters are estimated in the frequency domain and used as features for texture boundary detection. They obtained reasonable results in tests using checker board images of Brodatz textures. Later, for classification and shape from texture purposes, Choe and Kashyap [Choe91a] [Choe91b] used both first and second order models. Their hierarchical approach performs a first level classification based on surface roughness. This uses a *rotation invariant* first order model in which the two directional fractal differencing parameters are lumped into a single omnidirectional measure. The second level uses a more complex model. It employs two additional directional parameters ω_1 & ω_2 to account for *pattern* (as opposed to *roughness*) and also embodies surface tilt, slant, and rotation parameters. A maximum likelihood estimate of these parameters is made given the subgroup of classes indicated by the level one classification. Feature measure means and variances for a test set of Brodatz textures are reported. They indicate that classification would be successful for the majority of the textures, and that feature measure estimates for rotated, tilted and slanted textures differ very little from those of the original textures. Note again that this "rotation invariant" test uses rotated images of texture; a more realistic test would be to use images of rotated texture.

4.3.4. Markov random fields

Markov random fields are a two-dimensional generalisation of Markov chains which are defined in terms of conditional properties. The conditional probabilities are the same throughout the chain (or field) and are dependent *only* on a variable's neighbourhood (the Markov assumption). The size and form of the neighbourhoods are defined by the order of the model. The first, second, third, and fourth order neighbourhoods of *x* are shown below, where the first order neighbourhood consists of the variables labelled with a "1", the second order consist of all "1"s and "2"s etc.

	4	3	4	
4	2	1	2	4
3	1	*x*	1	3
4	2	1	2	4
	4	3	4	

Figure 4.1 - Markov random field neighbourhoods

Hassner and Sklansky [Hassner80] adapted a Markov random field (MRF) simulation algorithm originally used for gas models. They used it to generate first order isotropic MRF textures which lacked realism because they were binary. They pointed out that the number of parameters needed rose roughly with the square of the number of grey-levels. They however, also suggested that parameter estimation of MRF models could be used for classification purposes. In the special case where the texture is Gaussian, an MRF model may be parsimoniously characterised by a linear model [Chellappa85a].

Kashyap and Chellappa investigated parameter estimation schemes of, and texture synthesis with, both autoregressive and Markovian models. They concluded that selection between the two model types should depend entirely upon the data being considered [Kashyap83]. Chellappa et al further used MRF models for texture classification [Chellappa85a] and coding [Chellappa85b]. They tested a fourth order MRF model based feature set (i.e. 12 linear equation parameters). The test set consisted of 64x64 samples of Brodatz textures and they reported classification accuracies of 99% and 93% for the two feature sets respectively.

Cohen et al [Cohen91b] have also developed classifiers based on MRF models. Their application involves the detection of fabric defects, and in their tests a sixth order MRF parameter estimation scheme detected test defects with 100% accuracy.

Cohen also developed a *rotation and scale invariant* classification scheme using Markovian models [Cohen91a]. They developed a Gaussian MRF model that incorporates scale and rotation parameters, and showed that they could obtain estimates of these parameters given the normal MRF model. The classification scheme therefore first obtains maximum likelihood estimates for the scale and rotation factors of a test texture for each training class; and second, it determines the likelihoods of the test texture belonging to each class, given these estimated scalings and rotations. In a test using

Brodatz images they obtained good estimates of orientation and scaling factors, and 100% accurate classification results.

4.4. Non-model-based features

This section briefly reviews co-occurrence matrices, grey-level sum and difference histograms, Laws' masks, frequency domain methods, and Gabor filters.

4.4.1. Grey-level co-occurrence and other related features

Haralick [Haralick73] developed a set of fourteen feature measures based on "grey-tone spatial-dependence matrices", commonly referred to as grey-level co-occurrence matrices (GLCM). These matrices are essentially two-dimensional histograms of the occurrence of pairs of grey-levels for a given displacement vector. Typical displacement vectors include (1,0), (0,1), (1,1), (1,-1), (2,0). He achieved classification success rates of between 82% and 89% for photomicrographs of sandstones, aerial photographs, and satellite imagery.

Many researchers [Weska76] [Conners80] [Zucker80] [Davis81b] [Unser86] [Castrec88] [duBuf90] [Lovell92] [Shang93] have used Haralick's co-occurrence based features. The most popular features include Contrast, Angular Second Moment, Correlation, Inverse Difference Moment, and Entropy, with small displacement vectors e.g. (1,0) and (0,1) [Conners80].

Zucker [Zucker80] used a χ^2 test of independence for co-occurrence feature selection; the assumption being that the pairs of pixels would be independent of one another if the distance vector did not coincide with the structure of the texture. Lovell et al [Lovell92] used heuristic rules in a segmentation algorithm that combined Laws, co-occurrence, fractal dimension, and other feature measures; to produce good results on a diverse set of images which included an underwater image of an ROV (remotely operated vehicle).

Shang and Brown [Shang93] employed principal components analysis (PCA) to reduce the dimensionality of their co-occurrence based feature space. This improves the efficiency of training and classification sessions using neural network. They reported good results with Brodatz and side-scan sonar test sets.

Cheaper but related alternatives to grey-level co-occurrence features are those based upon grey-level *differences* [Weska76] [Conners80] [Davis81a] & [Castrec88]. These features are computed from histograms of differences between pairs of pixels (the pairs again defined by a displacement vector). Weska et al and Conners et al reported that performances from both feature sets are comparable, which is not surprising given their similarities. Their similarities are reinforced by the fact that the successful co-occurrence feature "contrast" can be computed directly from grey-level difference data. Unser went one step further and computed both sum and difference histograms [Unser86], from which he was able to calculate exact equivalents to nine of Haralick's co-occurrence features and estimate the remaining five. He used a test set of Brodatz images to compare his features against Haralick's. The results were almost identical.

Davis et al, generalised the grey-level co-occurrence matrix (GLCM) to take account of *any* features that may be generated from a pixel's neighbourhood (e.g. edge value and direction) rather than just their grey-levels [Davis81a]. The results of their experiments with a database of eight textures showed that the grey-level based descriptors (i.e. GLCM descriptors) gave the best results. Davis [Davis81b] also used "polarograms" as a way of showing the directional distribution of GLCM features, and proposed a set of polarogram statistics which are rotationally invariant. Haralick [Haralick73] had earlier suggested that *rotation invariant* features could be obtained from co-occurrence matrices by taking the average and range of each feature type over the four angles that he used. Rather than average existing features, Sun and Wee [Sun83] developed a directionally insensitive measure : the "neighbouring grey level dependence matrix". These matrices (indexed by grey-level k, and number of neighbours s) indicate the number of times a pixel of grey-level k, has s neighbours of a grey-level differing by less than a from k. Neighbourhoods are defined as all pixels within a specified radius. Hence the matrices, and features derived from them, contain no directional information. Despite this, tests using two features on Landsat images achieved percentage classification rates in the lower eighties.

4.4.2. Laws' texture energy filters

Laws [Laws79] [Laws80] investigated three texture feature generation methods in detail : co-occurrence, correlation, and spatial-statistical techniques. From the myriad of spatial statistical texture measures — essentially a set of statistical moments of a very wide and often ad hoc set of masks — and a desire to produce a computationally efficient method; Laws developed a coherent set of "texture energy" masks. All the masks were derived from three simple one-dimensional non-recursive filters. These may be convolved with each other to give a variety of one and two-dimensional filters. Laws found the most useful to be a set of seven bandpass and highpass directional filters, implemented as 5x5 masks. The outputs from these masks are passed to "texture energy" filters. These consist of a moving window calculation of variance (hence justifying the term "energy filter") or, more cheaply, a moving window average of absolute values. Laws used 15x15 windows, as a compromise between classification accuracy and computational cost. Texture energy images are used either directly, or via principal components analysis, as feature images for segmentation and/or classification.

Laws used Brodatz textures and other images to compare his masks with co-occurrence and correlation based features. He achieved pixel classification success rates of 94%, 72%, and 65% respectively. Castrec [Castrec88] however, found grey-level sum and difference based features to be superior for segmentation of side scan sonar images. Pietikainen et al [Pietikainen82] [Pietikainen83] tested Laws, co-occurrence contrast, and "edge per unit area" operators on Brodatz and geological terrain types. They found that the Laws operators performed consistently better than either edge or co-occurrence based features. Miller & Astley [Miller91] used Laws masks and morphological operators to detect glandular tissue in breast X-rays. They found that Laws masks R5R5, L5L5, and S5R5[1]; combined with a 31x31 local variance filter gave good results.

Greenhill and Davies [Greenhill93] employed 3x3 Laws' masks in conjunction with a neural network classifier. The output of the neural net is passed through a mode filter to

[1]Laws' masks are defined in chapter 5.

remove small areas which have been incorrectly classified. They reported the results of a set of experiments on a Brodatz montage that used a variety of sizes of averaging and mode filters. They concluded that the optimum sizes for these two filters are 11x11 and 13x13 respectively, and that mode filters represent a valuable but underused technique.

Harwood et al [Harwood85] reported 92% and 94% success rates for classification of 120x120 samples of six Brodatz textures using L5E5 and L5S5 masks respectively. DuBuf et al [duBuf90] used the variances of the masks' outputs within a relatively small 7x7 window and concluded that Laws' features were among the best tested out of a wide variety of texture measures.

In summary therefore a number of researchers in addition to Laws himself have found these easily implementable "texture energy measures" to compare very well with alternative approaches.

4.4.3. Frequency domain methods

Two-dimensional power or magnitude spectra provide information on texture coarseness and directionality from their radial and angular distributions respectively [Weska76]. The most commonly extracted features consist of sums of coefficients within *wedges, rings,* or *sectors* of two-dimensional power spectra [Lendaris70] [Kruger74] [Weska76] [He88]. Weska et al [Weska76] and Conners and Harlow [Conners80] found these features to give inferior results to those obtained using grey-level co-occurrence or difference based features. Other frequency domain measures include those derived from the characteristics of "spectral peaks". D'Astous and Jernigan [D'Astous84] used features which included the frequency (f), direction (θ), area, and relative power of spectral peaks. They tested their features against co-occurrence matrices and concluded that the former provided a higher level of discrimination between a test set of Brodatz textures. He et al [He88] later used the same spectral peak features together with co-occurrence and power spectra sector based measures on a Brodatz test set. However, when they used stepwise forward feature selection to select ten texture measures, they found six of the first seven were co-occurrence based and the remainder were derived from power spectra sectors, i.e. none of the first ten were derived from spectral peaks. The disappointing results of such frequency based techniques are not surprising given the difficulty of obtaining reliable spectral estimates from small samples of random signals [Marple87].

4.4.4. Gabor filters

Related to the Fourier based techniques described previously, are those that use banks of filters to highlight sections of two-dimensional spectra. Unlike the pure Fourier techniques however, the output is a set of images (one for each filter) that retain spatial information, and can therefore be used for segmentation purposes. Gabor filters are popular because the human vision system is also thought to employ similar banks of directional bandpass filters [Jain91]. Bovik et al provided an accessible description of Gabor filters in [Bovik87] where they also describe the separate use of magnitude and phase outputs for segmentation. Jain and Farrokhnia [Jain91] used banks of Gabor filters, followed by energy filters similar to those used by Laws [Laws80]. Segmentation tests on a variety of texture combinations, including Brodatz and MRF textures, gave good results, with *pixel* miss-classification rates ranging between 0.5% an 13%.

4.5. Comparative studies

This section collects together various comparative studies that have been mentioned in the above review. Its purpose is to allow a comparison to be made from "independent" tests of feature measures. Unfortunately these comparative studies do not cover all the features described above and in particular model-based features are poorly represented.

Weska et al [Weska76] investigated three types of texture measure for the classification of aerial and Landsat images. They used features based upon angular and radial power spectra, grey-level co-occurrence, grey-level difference, and grey-level run length. From the results of their experiments they concluded that

(i) Features based on grey-level difference or co-occurrence measures gave similar performances. Both gave better results than features derived either from power spectra or run length statistics.

(ii) The computationally cheapest feature, the mean of the grey-level differences, did about as well as other grey-level difference and co-occurrence measures.

Conners and Harlow [Conners80] reported a theoretical comparison of four types of texture measure that Weska et al investigated empirically. They measured the "amount of texture-context information contained in the intermediate matrices of each algorithm", using a set of synthetically generated Markovian textures. Their conclusions were similar to those of Weska et al.

Du Buf, Kardan & Spann [duBuf90] tested the ability of seven types of feature measure (computed in a 7x7 mask) to segment a set of synthetic test textures. They concluded that co-occurrence and Laws gave among the best results. The grey-level co-occurrence features were calculated using images requantised to 64 grey-levels. Distance vectors of (0,1) and (1,0) and features *contrast* and *difference variance* were found to give good results for single feature segmentation, while a combination of (1,0), (0,1), (1,1) and (1,-1) directions for the *contrast* feature gave the best multi-feature segmentation results. Of the Laws operators R5R5, E5L5, E5S5, and L5S5, managed to segment most of the test images. It is interesting to note that they describe these 5x5 masks as low cost Gabor functions. They also used fractal dimension but obtained "disappointing" results which they attributed to the coarse estimation methods they employed.

Castrec and Kernin [Castrec88] investigated the use of grey-level co-occurrence matrices, grey-level sum and difference histograms, and Laws' texture energy filter features. They applied these measures to the task of side scan sonar image segmentation. They found that the Laws features did not perform as well as the other two techniques, and that the co-occurrence contrast, correlation and variance features could be calculated 60 times more efficiently from grey-level sum and difference histograms. Their conclusion therefore was that features based on the latter technique were the most promising.

He et al [He88] did not perform an explicit comparison of texture feature performance, rather they used a standard feature selection procedure (forward stepwise) to select ten feature measures out of a mix of co-occurrence, PSD, and spectral peak based texture measures. Six out of the first seven selected were derived from co-occurrence matrices, while the rest were PSD sectors. None were based upon spectral peaks. This confirms the results reported by Weska et al, and Conners and Harlow, that Fourier spectrum based features do not seem to perform as well as their co-occurrence based competitors.

4.5.1. A league table of feature measures

The table below summarises the findings of comparative studies described above. The numbers indicate the relative order, in terms of classification and/or segmentation performance, that each study placed the particular feature concerned. Note that there are several *joint* placings and that 1 = first place (i.e. best). Blanks indicate that the feature was not investigated by the researchers.

	Co-occurrence	Sum and difference	Laws' masks	Run length	Fractal dimension	PSD wedges and rings	PSD peaks
Weska	1	1		2		2	
Conners	1	1		2		2	
duBuf	1		1		2		
Castrec	1	1	2				
He	1					2	3

Table 4.1 - Comparative studies of texture features

As the educational institutions are well aware, league tables should be treated with caution. However, from the above two points are immediately apparent. Comparative studies have tended to concentrate on non-model-based approaches, and of the features tested, co-occurrence matrices and their cheaper alternatives, the sum and difference based features, are reportedly better than other approaches.

4.6. Rotation invariance

A number of the papers discussed above have reported the development of *rotation invariant* texture features. It might be expected that such research would encompass the effects of variation in illuminant vector. However, few of the papers reviewed in this chapter discuss this topic and none investigate it in detail. They therefore implicitly assume that

(i) the textures concerned consist of surface markings only (albedo texture), or that

(ii) the illuminate vector is perpendicular to the surface of the texture, or that

(iii) the illumination is omnidirectional (flat).

The use of Brodatz textures as test cases reinforces the above assertion, as the only practical way of obtaining *rotated* examples of these textures is to rotate Brodatz's album before scanning, or alternatively to rotate the images once they have been scanned in. In either case this is *clearly not the same as rotating the physical textures themselves,* as the illumination is effectively rotated with the texture. Note that many of the Brodatz textures were photographed using directional-lighting to highlight surface relief.

For ease of reference the "rotation invariant" feature measures discussed above will now be summarised. They naturally fall into two camps. Firstly there are those that ignore directional information completely, i.e. they only consider omnidirectional measures or averages of directional features. Secondly there are those that exploit *relative* directional information, i.e. directional information that is independent of absolute angle, such as the angle *between* the two major directions in the texture etc.

4.6.1. Omnidirectional feature measures

Haralick [Haralick73] suggested computing omnidirectional features from directional measures by averaging his co-occurrence based features over the four directions. Sun and Wee [Sun83] took this to its logical conclusion and computed omnidirectional features directly from an omnidirectional matrix : the neighbouring grey-level dependence matrix whose entries depend on the values of neighbours in all directions. Kashyap and Khotanzad [Kashyap86] developed a *circular* autoregressive model which simply averaged all the pixels on the unit circle neighbourhood into a single value associated with a single parameter — producing a model containing no directional information. Mao and Jain [Mao92] took this one stage further by increasing the order of this rotation invariant autoregressive model to take in neighbourhood pixels on larger radii.

4.6.2. Rotation invariant directional feature measures

This section briefly describes texture features that measure directional characteristics of texture, yet are "rotation invariant". Such features include *relative* angular measures, such as the angle *between* the two major directions in the texture etc. Eichmann [Eichmann88] performed a Hough transform and then extracted rotation invariant features from parameter space. Although they were rotationally invariant they did exploit directional information : the angles between sets of lines, the spacing between parallel lines, and the number of principal line directions. Davis [Davis81b] constructed "polarograms" from GLCM features, from which he computed rotation invariant moments which he used as feature measures. Kashyap and Khotanzad [Kashyap86] augmented the omnidirectional "circular" autoregressive model with a directional one, from which they extracted a measure of directionality. Cohen et al [Cohen91a] integrated scaling and rotation parameters into an MRF model. Choe and Kashyap [Choe91a] [Choe91b] used a two level classification system : firstly an omnidirectional fractional differencing model was used for coarse classification according to texture *roughness*, and secondly, a second order directional model with rotation, tilt and slant parameters was used to refine the classification. Clarke [Clarke92] used principal components analysis on each texture region separately and then performed classification on the resulting principal components.

4.7. Conclusion

The preceding sections have reviewed a number of model-based and non-model-based feature measures, their purpose being to facilitate the selection of three types of texture measure for investigation as regards the effects of variation in illuminant vector. Papers dealing with illuminant vector effects are therefore particularly relevant to this selection. Unfortunately the author has not been able to find investigations of such effects in any of the literature reviewed above. This is particularly surprising in the case of rotation-invariant schemes.

The selection of texture measures for investigation has therefore used the following rather pragmatic criteria :

(i) popularity in the literature,

(ii) ease of implementation and use, and

(iii) performance.

From the table summarising the results of comparative studies the co-occurrence based features stand out as being prime candidates due to their popularity and performance. They were therefore selected for investigation. Texture measures based on Laws' masks were also selected. These features are particularly simple and efficient, are popular in the literature, and have the added bonus that they may be easily implemented in hardware. The third feature set was selected on a mixture of pragmatics and performance : Linnett's fractal inspired operator had been used by the author for segmentation of underwater images [Chantler91], and perhaps more importantly had achieved good results at the hands of Linnett and his colleagues e.g. [Linnett91a] [Clarke92] [Linnett93].

To summarise — the features selected for investigation as to the effects of illuminant vector variation are :

(i) Laws' texture energy masks [Laws80],

(ii) co-occurrence features [Haralick73], and

(iii) Linnett's fractal based operator [Linnett91a].

Chapter 5

Texture features and illumination

Chapters 2 and 3 of this thesis have shown that image texture is affected by changes in lighting and proposed that normalisation may compensate for slant angle variation. However, for isotropic texture variation in illuminant tilt introduces changes in the directional characteristics of the image which may not be compensated for in the same manner. Chapter 4 reviewed texture measures and selected three sets of features for further investigation as regards illumination effects. This survey also showed that little had been published on the effects of illuminant variation on texture classification. Hence the main purpose of this chapter is to determine the effects of changes in the illuminant's tilt and slant on the three feature sets.

The feature sets chosen in the preceding chapter for further investigation are : (i) Laws' masks, (ii) co-occurrence features, and (iii) Linnett's operator. Thus this chapter comprises three main sections — one for each feature set. Each of these sections is further sub-divided to address three aspects of feature set behaviour. Firstly, as the image model in chapter 2 and the subsequent empirical investigation in chapter 3 were based in the frequency domain, the frequency responses of the features is examined. This both provides a common view of their directional characteristics and gives an insight into their tilt and slant angle responses. Secondly the tilt and slant angle responses of the features applied to images of isotropic and directional texture are presented; and thirdly the effect of normalisation is investigated.

5.1. Laws' masks

Laws [Laws79] [Laws80] developed a set of two-dimensional masks derived from three simple one-dimensional filters.

They are :

L3 = (1,2,1) - Level detection,

E3 = (-1,0,1) - Edge detection, and

S3 = (-1,2,-1) - Spot detection.

Laws convolved these with each other, to provide a set of symmetric and anti-symmetric centre-weighted masks with all but the level filters being zero sum. These were convolved in turn with transposes of each other to give various sizes of square masks. He found the most useful to be those shown below. Note that the letters used in the mnemonics stand for Level, Edge, Spot, and Ripple.

-1	-2	0	2	1
-4	-8	0	8	4
-6	-12	0	12	6
-4	-8	0	8	4
-1	-2	0	2	1

L5E5

-1	0	2	0	-1
-2	0	4	0	-2
0	0	0	0	0
2	0	-4	0	2
1	0	-2	0	1

E5S5

1	-4	6	-4	1
-4	16	-24	16	-4
6	-24	36	-24	6
-4	16	-24	16	-4
1	-4	6	-4	1

R5R5

-1	0	2	0	-1
-4	0	8	0	-4
-6	0	12	0	-6
-4	0	8	0	-4
-1	0	2	0	-1

L5S5

Figure 5.1 - Four of Laws most successful masks (note the above would normally be used in conjunction with E5L5, S5E5, & S5L5 : the transposes of L5E5, E5S5, & L5S5)

The above masks are convolved with the original image to produce a number of images which are themselves passed through a second stage, which Laws termed a "macro statistic" [Laws79]. This consists of a moving window estimation of the energy within the images. Thus Laws' feature measures estimate the energy within the passband of their associated filters and he therefore called his operators "texture energy measures". He noted that variance is defined in terms of a sum of squares partly for mathematical convenience and proposed as an alternative, a cheaper but approximate measure : the average of the absolute values (ABSAVE). He found this to be just as successful, and as it requires less computation it will normally be used here.

As the masks are made up by convolving two one-dimensional components they are separable [Lim90], that is :

$$H(\omega_1,\omega_2) = H_1(\omega_1)H_2(\omega_2) \tag{5.1}$$

where

$H(\omega_1,\omega_2)$ is the frequency response of the two-dimensional mask,

$H_1(\omega_1)$ and $H_2(\omega_2)$ are the frequency responses in the x and y directions respectively, and

ω_1 and ω_2 are the angular frequencies in the x and y directions respectively.

Hence the frequency responses of the one-dimensional filters will be presented as a precursor to a description of the two-dimensional cases. The latter provide insight into the directionality of the operators and their response to image texture; a frequency domain model of which was presented in chapters 2 and 3. These frequency responses are followed by an examination of the effects of illuminant variation, using both the previously developed image model, and empirical observations. The issue of normalisation is also addressed.

5.1.1. Frequency response

a) One-dimensional frequency responses

The seven two-dimensional masks above may be obtained from four one-dimensional non-recursive filters, the weights of which are defined below :

L5 = (1,4,6,4,1)

E5 = (-1,-2,0,2,1)

S5 = (-1,0,2,0,-1)

R5 = (1,-4,6,-4,1)

The magnitude frequency response of L5 is simply obtained :

$$\begin{aligned} \left|H_{L5}(\omega_1)\right| &= \left|e^{-j2\omega_1} + 4e^{-j\omega_1} + 6 + 4e^{j\omega_1} + e^{j2\omega_1}\right| \\ &= 4(1+\cos\omega_1)^2 \end{aligned} \tag{5.2}$$

Note that as L5 = L3*L3 (where * represents the convolution operator) its magnitude response may be obtained from that of the L3 filter $\left|H_{L3}(\omega_1)\right| = 2(1+\cos\omega_1)$. Similarly for E5, S5 and R5. Thus :

$$\begin{aligned}\left|H_{E5}(\omega_1)\right| &= \left|H_{L3}(\omega_1).\,H_{E3}(\omega_1)\right| \\ &= 4\sin\omega_1\left(1+\cos\omega_1\right)\end{aligned} \tag{5.3}$$

$$\begin{aligned}\left|H_{S5}(\omega_1)\right| &= \left|H_{E3}(\omega_1).\,H_{E3}(\omega_1)\right| \\ &= 4\sin^2\omega_1\end{aligned} \tag{5.4}$$

$$\begin{aligned}\left|H_{R5}(\omega_1)\right| &= \left|H_{S3}(\omega_1).\,H_{S3}(\omega_1)\right| \\ &= 4\left(1-\cos\omega_1\right)^2\end{aligned} \tag{5.5}$$

L5, S5, and R5, are zero phase lowpass, bandpass, and highpass filters. E5 is a bandpass filter which introduces a phase change of 90° and whose passband is between those of L5 and S5. This is confirmed by figure 5.2, which contains plots of theoretical and empirical responses of the above one-dimensional features. The empirical results were obtained by applying the features to synthetically generated sine wave images followed by processing with the ABSAVE macro statistic (average of the absolute values).

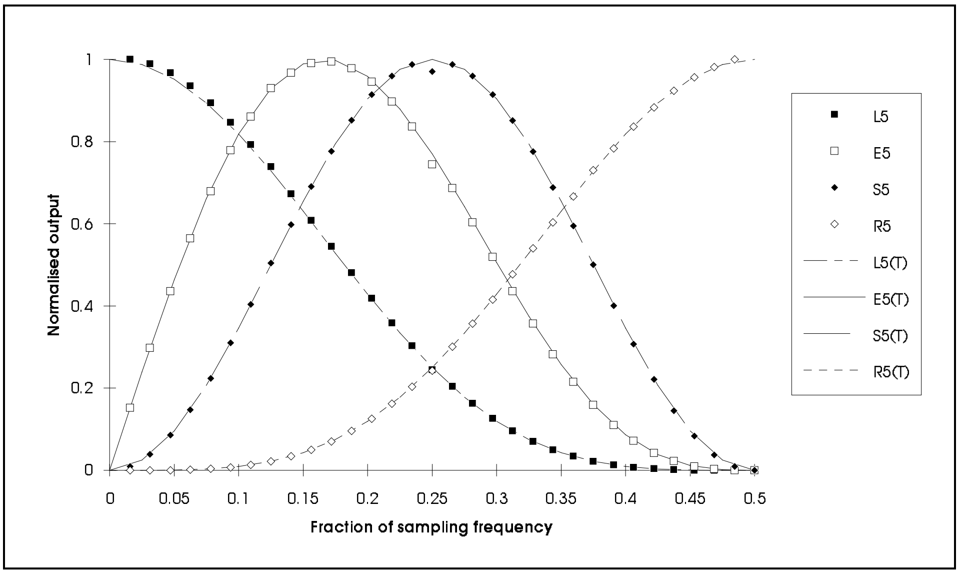

Figure 5.2 - Laws' one-dimensional operators : observed and theoretical (T) frequency responses

The observed responses of the one-dimensional feature measures match well with the theoretically derived results.

b) Two-dimensional frequency responses

Since Laws' masks are made up from separable one-dimensional filters, their frequency response may be simply obtained by substituting into (5.1), i.e. by multiplication in the frequency domain :

$$\begin{aligned}\left|H_{L5E5}(\omega_1,\omega_2)\right| &= \left|H_{L5}(\omega_2).\,H_{E5}(\omega_1)\right| \\ &= 4(1+\cos\omega_2)^2\,4\sin\omega_1\left(1+\cos\omega_1\right)\end{aligned} \tag{5.6}$$

$$\begin{aligned}\left|H_{E5S5}(\omega_1,\omega_2)\right| &= \left|H_{E5}(\omega_2).\,H_{S3}(\omega_1)\right| \\ &= 4\sin\omega_2\left(1+\cos\omega_2\right)4\sin^2\omega_1\end{aligned} \tag{5.7}$$

$$\begin{aligned}\left|H_{R5R5}(\omega_1,\omega_2)\right| &= \left|H_{R5}(\omega_2).\,H_{R5}(\omega_1)\right| \\ &= 4\left(1+\cos\omega_2\right)^2 4\left(1+\cos\omega_1\right)^2\end{aligned} \tag{5.8}$$

$$\begin{aligned}\left|H_{L5S5}(\omega_1,\omega_2)\right| &= \left|H_{L5}(\omega_2).\,H_{S5}(\omega_1)\right| \\ &= 4\left(1-\cos\omega_2\right)^2 4\sin^2\omega_1\end{aligned} \tag{5.9}$$

and for the relevant transposes :

$$\begin{aligned}\left|H_{E5L5}(\omega_1,\omega_2)\right| &= \left|H_{E5}(\omega_2).\,H_{L5}(\omega_1)\right| \\ &= 4\sin\omega_2\left(1+\cos\omega_2\right)\ 4(1+\cos\omega_1)^2\end{aligned} \tag{5.10}$$

$$\begin{aligned}\left|H_{S5E5}(\omega_1,\omega_2)\right| &= \left|H_{S5}(\omega_2).\,H_{E5}(\omega_1)\right| \\ &= 4\sin^2\omega_2\,4\sin\omega_1\left(1+\cos\omega_1\right)\end{aligned} \tag{5.11}$$

$$\begin{aligned}\left|H_{S5L5}(\omega_1,\omega_2)\right| &= \left|H_{S5}(\omega_2).\,H_{L5}(\omega_1)\right| \\ &= 4\sin^2\omega_2\,4\left(1-\cos\omega_1\right)^2\end{aligned} \tag{5.12}$$

In addition to the above theoretically derived responses, empirical results were also obtained. Sets of "corrugated" sine wave surfaces were used as inputs to the feature measures and the average output measured. The results are shown below. Since the empirical plots were similar to the theoretical responses only the former are shown.

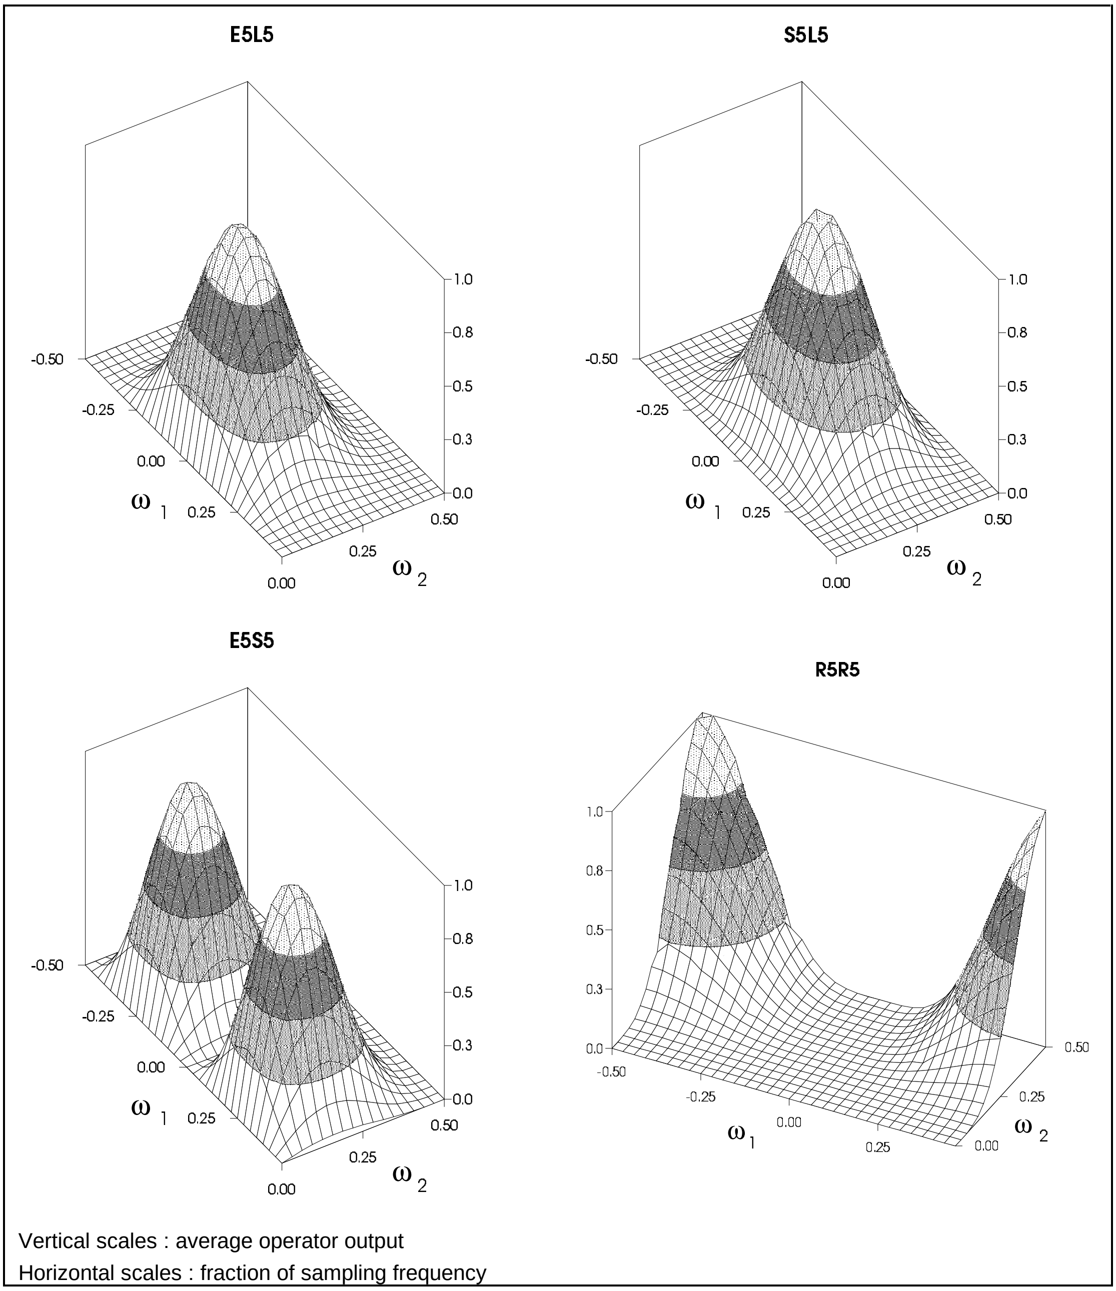

Figure 5.3 - Laws' operators : empirical two-dimensional frequency responses

The above graphs show that E5L5 and S5L5 (and hence their transposes) are uni-directional, while E5S5 is bi-directional. What is interesting however, is that the mask of the R5R5 feature which at first glance appears to be isotropic is in fact bi-directional; being sensitive to high frequencies at 45° and 135°. Thus, with the exception of L5L5, all of Laws' masks are directional and all are likely to be affected by variation in illuminant tilt.

5.1.2. Tilt angle response

This section investigates the response of Laws' operators to changes in the tilt angle of the illuminant. Firstly the theoretical tilt response of the L5E5 uni-directional operator is examined using the image model of topological texture developed previously, and the operator's theoretical frequency response. The resulting predictions are compared with empirical results obtained from laboratory experiments. Secondly, the empirical response of Laws' bi-directional operators to isotropic and directional textures is presented. Thirdly, the effects of image normalisation are investigated, with the aim of assessing whether or not such a procedure compensates for the effects of variation in tilt.

a) The tilt response of the uni-directional operator L5E5

This section examines the theoretical response of laws' L5E5 operator; which is obtained from the product of its transfer function and the frequency domain model of image texture developed in chapters 2 and 3. These results are compared with those obtained from laboratory experiment. The purpose of this investigation is two-fold : firstly it is to establish the tilt response of the operator and secondly it is to show the utility of the image model developed earlier.

As only variations due to changes in the illuminant's tilt are of interest, it is assumed that the illuminant's slant does not vary, and the contribution of the corresponding component in the image model is a constant k_σ. Thus the model presented in equations (2.14 to 2.17) reduces to :

$$F_I(\omega,\theta) = F_s(\omega,\theta).F_\tau(\omega,\theta).k_\sigma \tag{5.13}$$

Now if the test textures are assumed to be isotropic and the radial shapes of the log-log magnitude spectra assumed to be straight lines, then the magnitude of the surface response component may be represented by :

$$\left|F_s(\omega,\theta)\right| = k_\beta \omega^{-\beta_I/2} \tag{5.14}$$

The parameters k_β and β_I may be estimated by obtaining the gradient and y-intercept of the best-fit straight line to the average log-log radial plots. Furthermore in chapter 3 the directional characteristics of these textures was shown to approximate to a raised cosine (3.8) i.e.

$$F_\tau(\theta) = (m_\tau \cos(\theta - \tau) + b_\tau)^1 \tag{5.15}$$

and the parameters m_τ and b_τ were estimated for each of the test textures (see table 3.2). In order to prevent estimations of the power of the spectra being included twice in the model the directional characteristics were modelled by a *normalised* tilt angle component $F'_\tau(\theta)$.

$$F'_\tau(\theta) = (m'_\tau \cos(\theta - \tau) + b'_\tau) \tag{5.16}$$

where

$$m'_\tau = {}^{m_\tau}\!/_{F_\tau(\theta)} \quad \text{and} \quad b'_\tau = {}^{b_\tau}\!/_{F_\tau(\theta)}$$

Hence the image texture magnitude spectra of the four samples may be modelled by combining (5.13), (5.14) and (5.16) to give :

$$|F_I(\omega,\theta)| = k_\beta \omega^{-\beta_I/2} \left(m'_\tau \cos(\theta - \tau) + b'_\tau\right).k_\sigma \tag{5.17}$$

Now if only relative magnitudes are required, k_σ may be eliminated and all remaining parameters estimated for each of the test textures as described previously.

Thus the output of the first stage of Laws' operators is simply derived, e.g. for L5E5 combining (5.6) and (5.17) gives :

$$\begin{aligned} |Y_{L5E5}(\omega,\theta)| &= |H_{L5E5}(\omega,\theta)||F_I(\omega,\theta)| \\ &= \sin(\omega\cos\theta)\{1+\cos(\omega\cos\theta)\}\{1+\cos(\omega\sin\theta)\}^2. \\ &\quad k_\beta \omega^{-\beta/2}\{m'_\tau \cos(\theta-\tau) + b'_\tau\}.k_\sigma \end{aligned} \tag{5.18}$$

where :

$\omega\cos\theta = \omega_1$, $\omega\sin\theta = \omega_2$, and

$|Y_{L5E5}(\omega,\theta)|$ is the two-dimensional magnitude spectrum of the output of L5E5.

As previously discussed, the second stage of Laws' operators use variance or, more cheaply, the *average of absolutes* as an "energy measure". Normally the latter is used in

---

[1] ω has been omitted here as the *tilt component* is not a function of frequency.

this thesis. In this section however, the average of the squares will also be used (this is identical to the variance for zero-mean images). The latter is used in this section because it is more tractable analytically. It provides an estimate of the "power" of the image texture and hence the integral of the PSD [Ogilvy91] [Cooper86]. Note this assumes that the process under consideration is at least wide-sense stationary [Peebles87]. The PSD $S(\omega,\theta)$ may in turn be obtained from the magnitude of the Fourier transform of the output of the filter. Hence the mean output of the L5E5 operator will be :

$$
\begin{aligned}
\overline{y_{L5E5}^2} &= \frac{1}{4\pi^2}\int_{\omega=0}^{\infty}\int_{\theta=0}^{2\pi} S(\omega,\theta)\,d\theta d\omega \\
&= \frac{1}{4\pi^2}\int_{\omega=0}^{\infty}\int_{\theta=0}^{2\pi} \left|Y_{L5E5}(\omega,\theta)\right|^2 d\theta d\omega \\
&= \frac{1}{4\pi^2}\int_{\omega=0}^{\infty}\int_{\theta=0}^{2\pi} \left[\sin(\omega\cos\theta)\{1+\cos(\omega\cos\theta)\}\{1+\cos(\omega\sin\theta)\}^2 . k_\beta \omega^{-\beta/2}\{m'_\tau\cos(\theta-\tau)+b'_\tau\}.k_\sigma\right]^2 d\theta d\omega
\end{aligned}
\tag{5.19}
$$

The solution of the above integral for the general case is not trivial. However, it may be estimated numerically when the values of the parameters are known. Hence the four parameters (m'_τ, b'_τ, k_β, and β_I) were calculated for each of the four isotropic test textures *beans1*, *chips1*, *rock1* and *stones1*; using the estimation techniques described previously. The integral (5.19) was evaluated for each of these four sets of parameter values for nineteen angles of illuminant tilt (0° to 180° in 10° steps). The results are shown in figure 5.4.

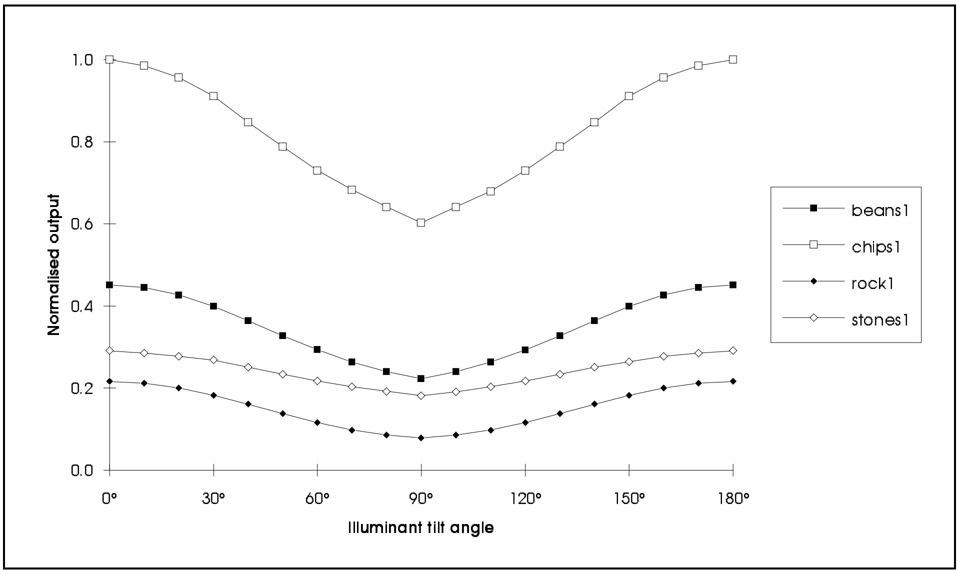

Figure 5.4 - Predicted effect of tilt angle variation on L5E5 output[2]

For comparison figure 5.5 shows the equivalent results obtained by processing images of the textures with the feature measure itself — an L5E5 mask coupled to a *mean square* macro statistic.

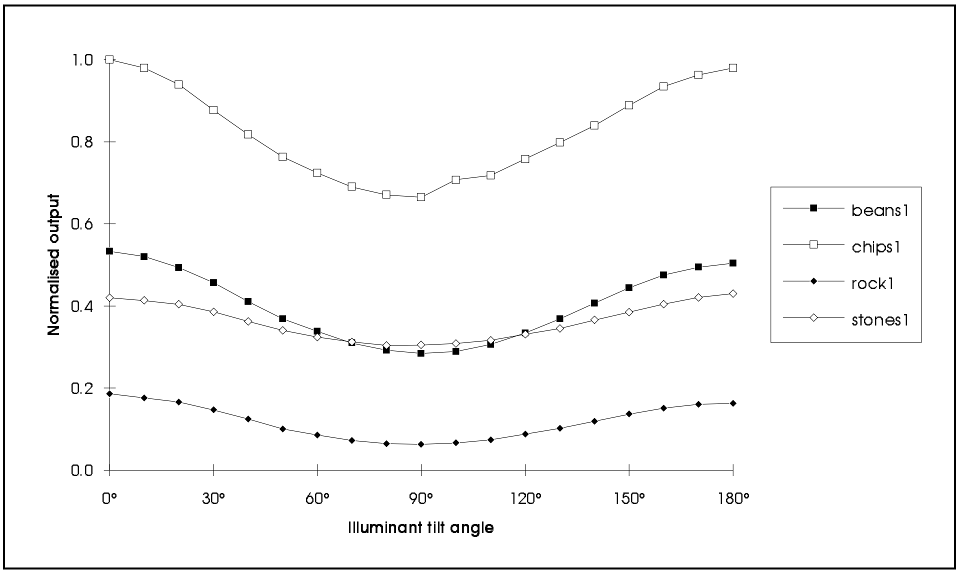

Figure 5.5 - Observed effect of tilt angle variation on L5E5 output[2]

The above graphs show that for the Laws' L5E5 operator, (i) the output is affected by variation in illuminant tilt, and (ii) the image model of the four isotropic textures developed in chapter 3 predicts the effects of variation in tilt reasonably well. For comparison the output of the same L5E5 mask but with the cheaper ABSAVE macro statistic is shown in figure 5.6. It can be seen that the cheaper *average of absolutes* macro statistic gives similar results to the *mean square* operator. Indeed for the case shown the former would seem to give better separation between the classes. Laws found little difference between the performance of these macro statistics and therefore preferred the cheaper ABSAVE operator. Hence this macro statistic will be used in the remainder of this document.

[2]Note that the data presented in figures 5.4 and 5.5 was obtained by taking averages of feature images, and that together with figure 5.3 these graphs have been scaled to a maximum value of 1.0 for comparison purposes.

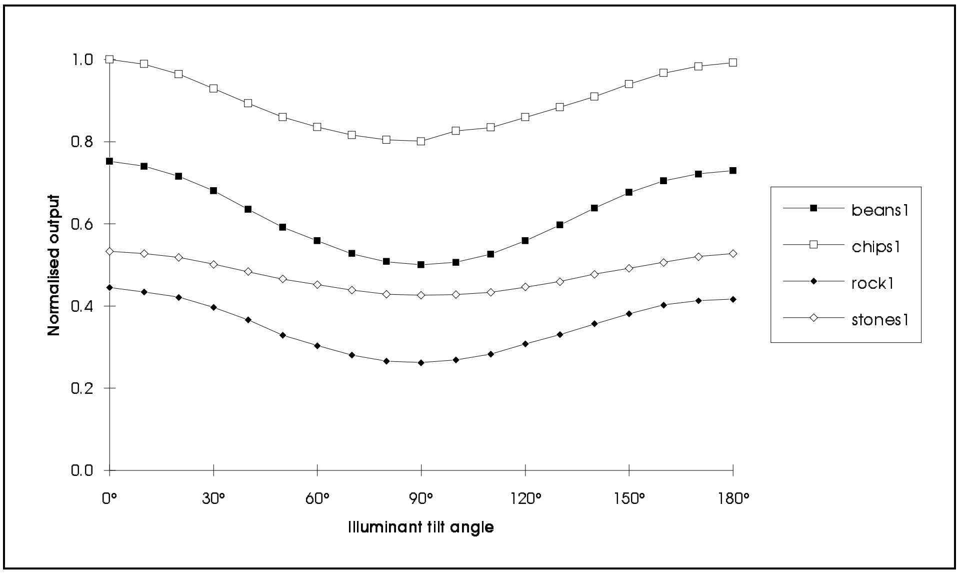

Figure 5.6 - Effect of tilt angle variation on L5E5 output (ABSAVE macro statistic)[2]

The behaviours of feature means are obviously important for classification and segmentation purposes, but they do not provide sufficient information to allow the likely effects to be assessed. What is required are the behaviours of the distributions. A small variation in mean due to change in illuminant tilt may be significant for distributions of large variance, but insignificant for those of small variance. For example the figure below shows distributions of L5E5 (with a 29x29 ABSAVE macro statistic) for two textures under two lighting conditions.

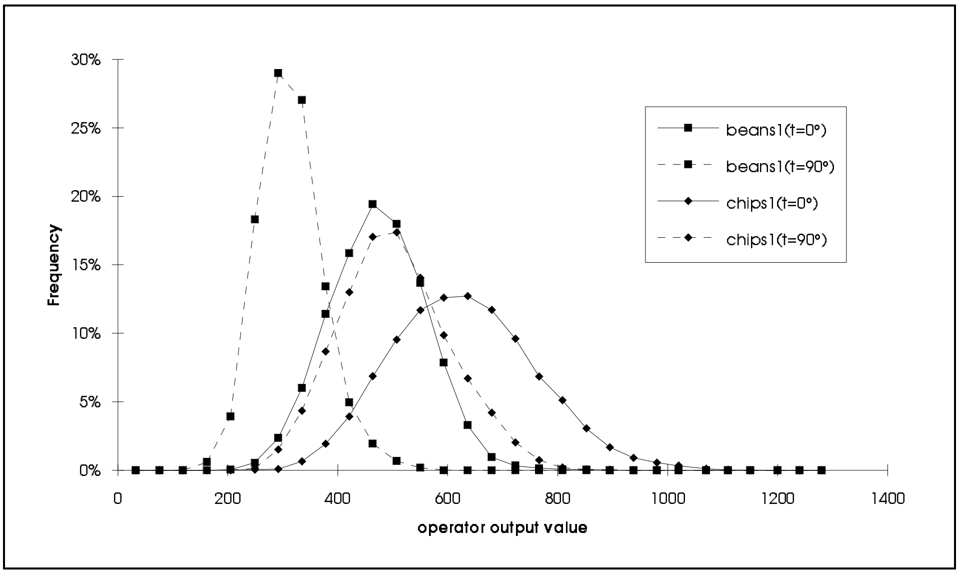

Figure 5.7 - Effect of tilt variation on L5E5 distributions

Assuming equal prior probabilities a maximum likelihood classifier trained under an illuminant tilt of 0° would have a decision surface at approximately L5E5 = 580. That is "halfway" between the *beans1* (τ = 0°) and *chips1* (τ = 0°) distributions (solid line graphs). However, the dashed graphs show the result of changing the tilt to 90° : the mean of *chips1* is now clearly to the left of L5E5 = 580, and so the majority of this class at τ = 90° would be mis-classified. Note that in this case increasing the window size of the macro statistic would be likely to *increase* the number of incorrectly classified *beans1* pixels — as it would most likely reduce the variance of this distribution.

Thus changes in the illuminant's tilt have been shown to affect the output of Laws' L5E5 operator. Experiments using the four isotropic test textures with the other uni-directional feature (L5S5) gave similar results to those shown above.

b) Bi-directional operators

So far only the behaviour of uni-directional operators has been considered, but what of the bi-directional operators ? Clearly illumination tilt will not affect isotropic operators when used on isotropic physical textures. However, R5R5 and other operators produced by convolving a one-dimensional bandpass or highpass filters with other similar filters are not isotropic. Instead they are bi-directional (being sensitive to diagonal or near diagonal components). Thus it would be reasonable to expect such directional filters to be affected by illuminant tilt. The figure below shows the tilt responses of two bi-directional operators obtained from four isotropic textures (*beans1, chips1, rock1,* and *stones1*).

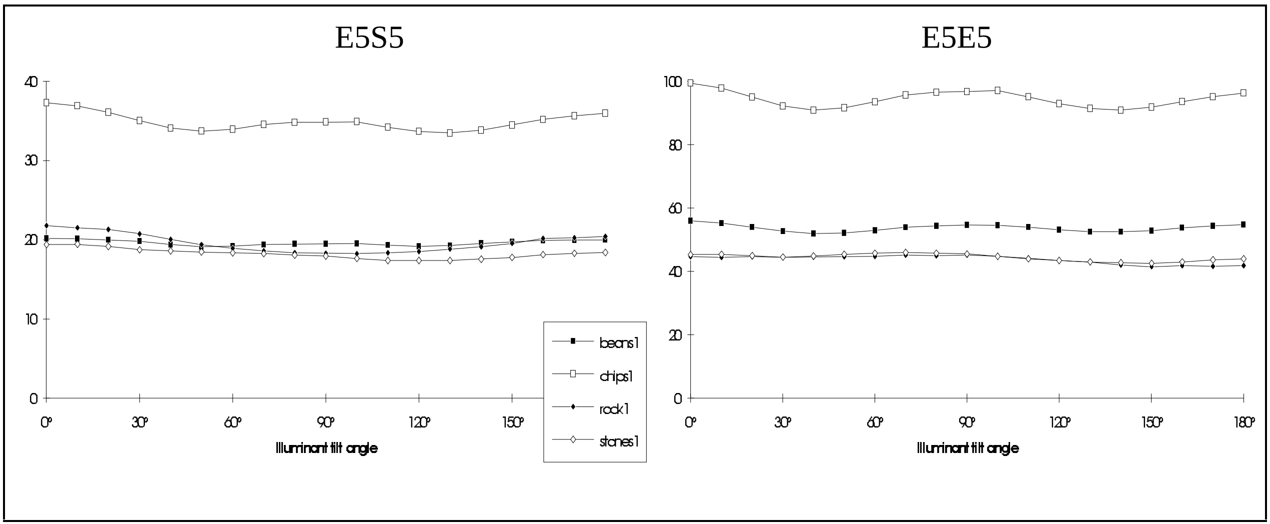

Figure 5.8 - Tilt response of the bi-directional operators E5S5 and E5E5

The E5S5 and E5E5 results show that although these operators are affected by tilt, the effects are not nearly as pronounced as for uni-directional operators. This may be explained by the fact that these feature measures are sensitive to two near mutually perpendicular directions, and as one is being attenuated by a particular illuminant tilt the other is being enhanced. Thus bi-directional operators with mutually perpendicular axes of sensitivity, will be least affected by illuminant tilt when used with isotropic textures. If the angle between the two axes is reduced, then the behaviour will tend towards that of the uni-directional case.

However, if the physical texture is not isotropic, then bi-directional features such as E5E5 may be significantly affected by illuminant tilt; as is shown in the figure below.

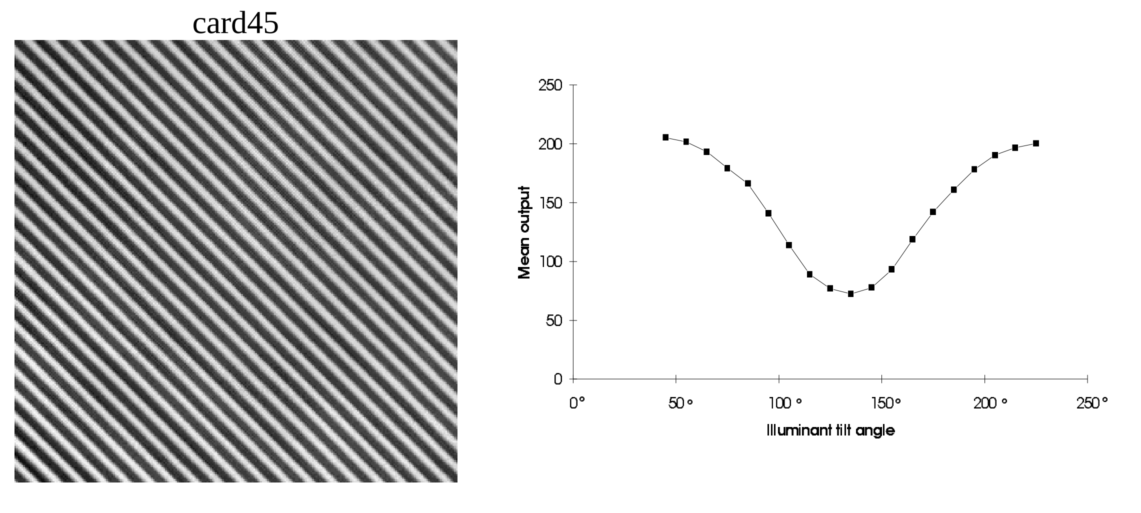

Figure 5.9 - The tilt response of E5E5 for the directional texture "card45"

The above shows a sample of uni-directional texture "*card45*", and the corresponding tilt response of E5E5. Here there is no compensating effect as was the case for the isotropic textures, and so the operator is significantly affected.

c) Normalisation

A number of texture classification schemes normalise image data in some manner to remove "brightness variation". This is usually performed either by histogram equalisation or by re-scaling the data to have a common mean and variance, e.g. see [Greenhill93] [Bovik87] [duBuf90] [Laws79] [Weska76] and [Haralick73]. Chapter 2's image model predicts that, for *isotropic* textures, normalisation could compensate for variation in illuminant slant but not tilt. Here therefore, the tilt response of the L5E5 operator applied to *normalised* image data is examined. Each image of the test set was scaled to have a mean of 127 and a variance of 100 (note that local brightness variation is compensated for by using registration images as described in chapter 3). The figure below shows the mean output of L5E5 for tilts of between 0° and 180° using normalised images of the four isotropic test textures.

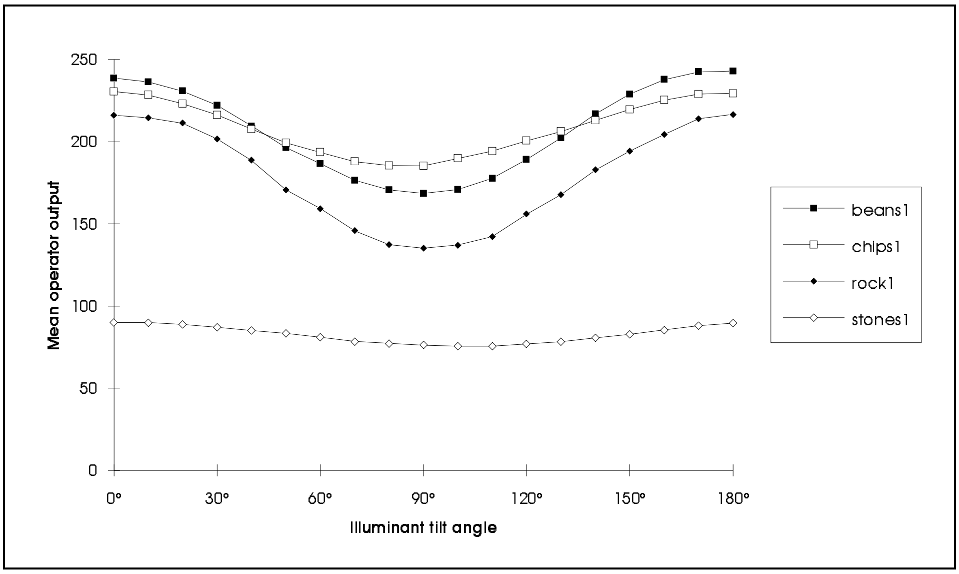

Figure 5.10 - The effect of normalisation on L5E5 tilt angle response

The above shows that normalisation does affect the tilt response but it certainly does not compensate for it. Indeed normalisation of the images has actually reduced the separation between the classes *beans1*, *chips1*, and *rock1*, thereby complicating the classification task for this operator (compare the above with figure 5.6). The closeness of the distributions is more clearly shown in the figure below.

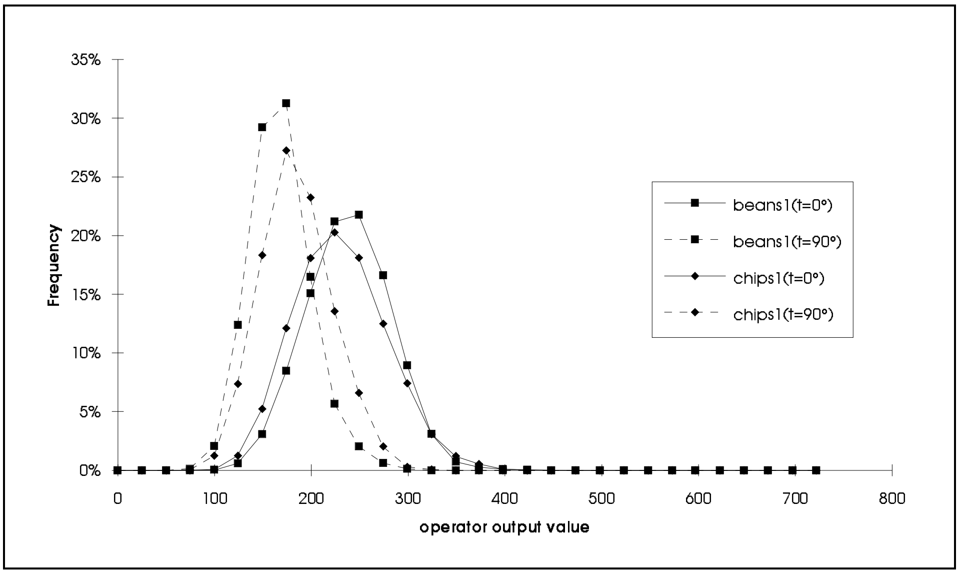

Figure 5.11 - Effect of normalisation on the L5E5 output distribution for tilts of 0° and 90°.

The solid line plots of figure 5.11 show that normalisation has made the distributions of L5E5, for the two textures *beans1* and *chips1,* almost identical for τ = 0°. However, variation of illuminant tilt (to τ = 90° - dashed line plots) still produces a significant change in mean values and considerable mis-classification would again occur. Note that normalisation will compensate for variation in the intensity of the illuminant, but that it also has the unfortunate effect of normalising a significant discriminatory feature : image variance.

5.1.3. Slant angle response

From chapter 2 the predicted illuminant slant angle (σ) response (2.17) is :

$$F_\sigma = \sin \sigma$$

As this is independent of frequency it effects a uniform amplification or attenuation of image texture across the spectrum. All of Laws' feature measures will therefore be affected in a similar manner as they provide an estimate of the power in their passbands. Un-normalised slant angle responses of L5E5, for the four isotropic test textures and a fifth uni-directional texture *card1,* are shown below. Note : *card1* is the corrugated cardboard surface shown in figure 5.9 except that the corrugations run vertically.

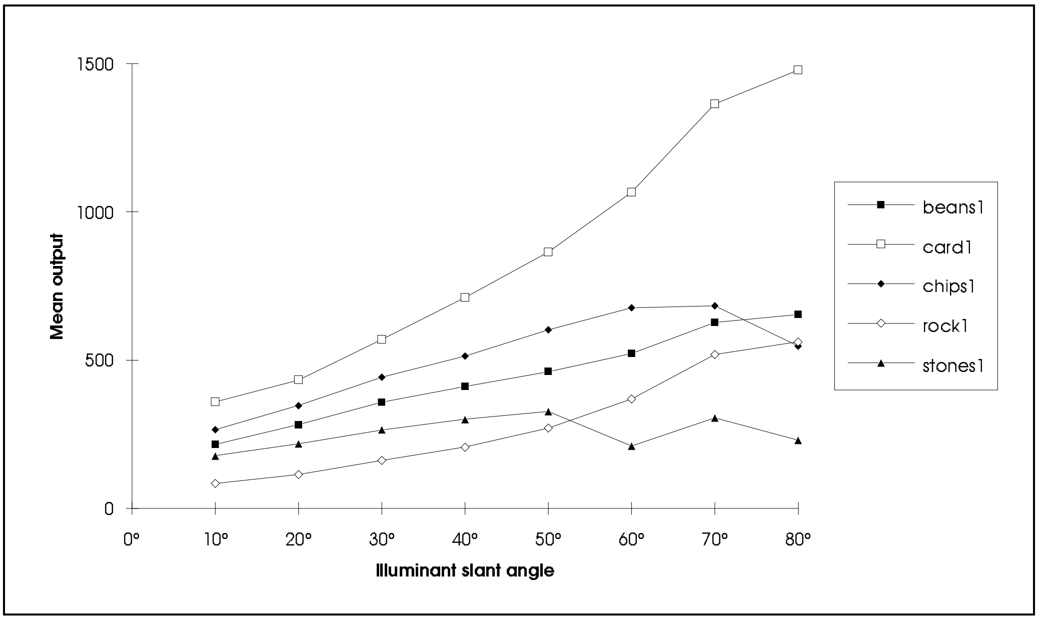

Figure 5.12 - Un-normalised slant angle response of Laws' L5E5 operator

The above figure shows that the L5E5 slant angle responses mimic those of the magnitude spectra of figure 3.35; that is there is a gradual increase in the mean output with increasing σ for all textures up until 50°, after which all but *stones1* and *chips1* continue to increase. Since shadows are longer and cover larger areas at higher slant angles, the power of the frequency components may decrease as the slant angle increases. Note that shadowing is particularly noticeable in the images of *stones1* and *chips1;* and that, for these two textures, the output of the L5E5 operator is reduced at higher slant angles.

What is important however, is that the illuminant slant angle does significantly affect the L5E5 operator when used on un-normalised images.

If change in illuminant slant does effect a uniform amplification/attenuation across the spectrum as suggested in chapter 2, then normalisation will compensate for these variations. In order to investigate this effect the test image sets were normalised as before to have a mean = 127 and a variance = 100. The Laws' L5E5 operator was applied to the resulting images. Its mean output, as a function of illuminant slant, is shown below.

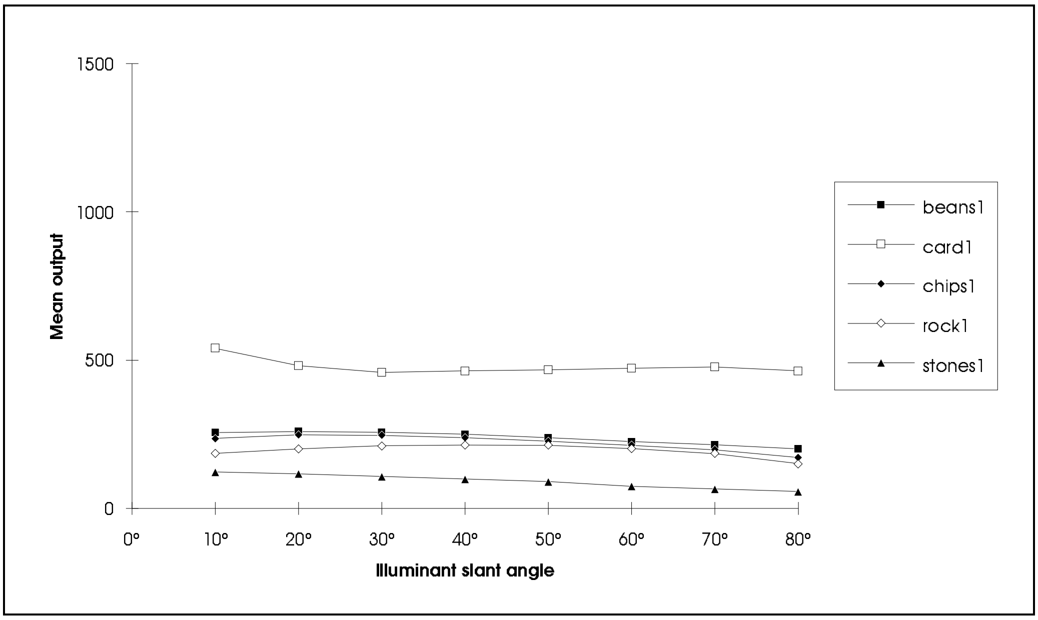

Figure 5.13 - Normalised response of Laws' L5E5 operator

When compared with the previous figure, the graphs above show that normalisation has significantly reduced the variation due to changes in slant. These results suggest that normalisation may reduce the effect of changes in illuminant slant on classification. However, this reduction in variation with slant angle has been bought at the expense of reduced separation of class means. Thus normalisation may actually increase classification errors rather than decrease them.

The subject of normalisation is further addressed in sections 5.2 and 5.3 on Linnett's and co-occurrence feature measures.

5.1.4. Summary

This section has investigated the response of Laws' operators to changes in the illuminant's tilt and slant angles. The following points summarise its findings.

- The two-dimensional magnitude frequency responses of the popular Law's operators have shown that they are either uni-directional or bi-directional.
- The image model developed in chapters 2 and 3 was used to predict the tilt response of the L5E5 operator. The results were similar to those obtained empirically showing (i) the utility of the image model, and (ii) that the L5E5 operator is not invariant to changes in illuminant tilt.

- The bi-directional operator E5E5 was less affected by changes in tilt when applied to images of isotropic textures, but it was significantly affected when used on the uni-directional texture *card45*.
- Image normalisation did not compensate for these variations as far as isotropic textures were concerned.
- The L5E5 operator was significantly affected by changes in illuminant slant.
- Normalisation reduced the variations due to changes in slant, but also reduced the separation between test textures' means.

5.2. Co-occurrence matrices

Co-occurrence matrices have been widely used for texture classification [Haralick73] [Weska76] [Conners80] [Zucker80] [Davis81b] [Unser86] [Castrec88] [duBuf90] [Lovell92], but perhaps because of their computational cost they have been used less frequently for segmentation [duBuf90]. A co-occurrence matrix is a two dimensional histogram of pixel pairs defined by a displacement vector **d**. They are an estimate of the joint probability function of these pixel pairs. Haralick [Haralick73] defined 14 statistics to provide an economical way of describing these distributions, and it is these that are used as features for texture discrimination. Here only four of the most popular will be investigated. They are ASM (angular second moment), ENT (entropy), COR (correlation), and CON (contrast) :

$$ASM = \sum_{i=0}^{Ng-1}\sum_{j=0}^{Ng-1} p(i,j)^2 \tag{5.20}$$

$$ENT = \sum_{i=0}^{Ng-1}\sum_{j=0}^{Ng-1} -p(i,j)\log(p(i,j)) \tag{5.21}$$

$$COR = \frac{1}{\sigma_x \sigma_y}\left[\sum_{i=0}^{Ng-1}\sum_{j=0}^{Ng-1}(ij)\; p(i,j) - \mu_x \mu_y\right] \tag{5.22}$$

$$CON = \sum_{i=0}^{Ng-1}\sum_{j=0}^{Ng-1}(i-j)^2\; p(i,j) \tag{5.23}$$

where

$p(i,j) = P(i,j)/n$,

$P(i,j)$ is the (i,j)th element of the un-normalised co-occurrence matrix defined by a displacement vector **d** and window W,

n is the normalising constant $n = \sum_{i=0}^{Ng-1} \sum_{j=0}^{Ng-1} P(i,j)$,

Ng is the number of grey-levels, and

μ_x, μ_y, σ_x, and σ_y are the means and standard deviations of the marginal distributions.

In the forms above it is difficult to determine their frequency response and they are also expensive to compute. The next section therefore derives alternative expressions for two of the features. They are formulated directly in terms of image grey-levels, and do not use co-occurrence matrices. These alternative expressions are used as the basis for efficient moving window implementations. In the case of the contrast operator the alternative expression is also used to derive its frequency response, which is presented in the following section along with empirical results for the other operators. Finally the results of laboratory experiments on illuminant variation are presented and the issue of normalisation is investigated.

5.2.1. An alternative formulation

The contrast feature is commonly used in an alternative form expressed directly in terms of grey-levels :

$$\begin{aligned} CON &= \sum_{i=0}^{Ng-1} \sum_{j=0}^{Ng-1} (i-j)^2 \; p(i,j) \\ &= \frac{2}{n} \sum_{(i,j) \in D} (i-j)^2 \end{aligned} \tag{5.24}$$

where

D is the set of pixel pairs defined by the displacement vector **d** within a window W.

This form is amenable to efficient implementation and frequency domain analysis. To the author's knowledge no such equivalent expressions have been published for the other features; although Unser did derive more efficient and in some cases approximate alternatives using sum and difference histograms [Unser86]. The correlation feature however, may be formulated directly in terms of grey-levels and statistics of the marginal distributions. From (5.22)

$$COR = \frac{1}{\sigma_x \sigma_y}\left[\sum_{i=0}^{Ng-1}\sum_{j=0}^{Ng-1}(ij)\ p(i,j) - \mu_x \mu_y\right]$$
$$= \frac{1}{\sigma_x \sigma_y}\left[\frac{2}{n}\sum_{(i,j)\in D} ij - \mu_x \mu_y\right] \tag{5.25}$$

Similarly for the marginal distribution statistics :

$$\mu_x = \sum_{i=0}^{Ng-1} i.\,p_x(i)$$
$$= \sum_{i=0}^{Ng-1} \frac{i.\,P_x(i)}{n} \tag{5.26}$$
$$= \frac{1}{n}\sum_{i\in D_2} i$$

$$\sigma_x{}^2 = \sum_{i=0}^{Ng-1} (i-\mu_x)^2\ p_x(i)$$
$$= \frac{1}{n}\sum_{i\in D_2} (i-\mu_x)^2 \tag{5.27}$$

where

D_2 is the set of pixel values contained within the pixel pairs in D.

Note that because p(i,j) is symmetric about the leading diagonal

$$\mu_x = \mu_y \tag{5.28}$$

$$\sigma_x{}^2 = \sigma_y{}^2 \tag{5.29}$$

and the normalising constant may be calculated directly :

$$n = \sum_{i\in D_2} 1$$
$$= 2(w_x - d_x)(w_y - d_y) \tag{5.30}$$

where

w_x and w_y = size of the window W in x and y respectively, and

d_X and d_Y = Cartesian components of the displacement vector **d.**

Using (5.25) to (5.29) the correlation feature may be expressed as

$$COR = \frac{1}{\sigma^2}\left[\sum_{(i,j)\in D} ij - \left(\frac{1}{n}\sum_{i\in D_2} i\right)^2\right] \tag{5.31}$$

where

$$\sigma^2 = \sigma_x{}^2 = \sigma_y{}^2 = \frac{1}{n}\sum_{i\in D_2} i^2 - \left(\frac{1}{n}\sum_{i\in D_2} i\right)^2 \tag{5.32}$$

Thus (5.31) and (5.32) allow the correlation feature to be calculated directly from image data, without need to reference the co-occurrence matrix itself. Indeed only three running totals need to be maintained in an incremental algorithm. They are :

$$\sum_{i\in D_2} i, \quad \sum_{i\in D_2} i^2 \ \text{ and } \ \sum_{(i,j)\in D} ij$$

Thus all variables (apart from image data) may be kept in registers and the computation incurred in matrix address calculations may be avoided.

The above expressions for COR, CON, were used as the basis for moving-window feature measures. They incorporate an incremental update mechanism similar to the technique used by Huang et al to create a fast median filter [Huang78]. The window is initialised in the top left corner of the image and moved across the image pixel by pixel. Each time the window is moved, the last column is removed, a new first column added, and intermediate and feature values are updated accordingly. Thus considerable re-computation is avoided. Although the ENT and ASM features do not lend themselves to transformation into a grey-level based formulation, the incremental update method using a moving window may be applied. This involves the maintenance of $P(i,j)$ in an incremental manner, but again the saving in processing time is considerable. For a 35x35 moving window the incremental form involves a little over 70 operations per window position, whereas a "straight" implementation involves 1225.

For classification purposes feature sets need only be calculated for each unknown region — whereas segmentation requires feature values to be generated for every pixel in the image. Thus the use of the alternative implementations, described above, makes the use of co-occurrence features for *segmentation* of sequences of 512x512 images practical on a standard workstation within two to three minutes.

5.2.2. Frequency response

The image model of topological texture previously developed, and the ensuing empirical investigations used frequency domain representations. In order to gain an insight into the effects of illuminant variation on co-occurrence features, and to provide an alternative view of their directional characteristics, their frequency responses will now be investigated. These responses should however be viewed with caution; as the co-occurrence operators are non-linear.

The contrast operator CON is straightforward to analyse and will now be presented. The three other features are not, and so they will only be investigated empirically.

a) Contrast feature : frequency response

From the image based form of the CON feature (5.24) it can be seen that it is simply an edge operator followed by square and average functions. The latter two functions form an *energy measure* — as used in Laws' filters. For a displacement vector **d =** (1,0) the edge operator becomes a horizontal non-recursive filter with weights of (-1,1), hence the frequency response of the CON(1,0) *filter* is :

$$\begin{aligned}\left|H_{CON}(\omega_1,\omega_2)\right| &= \left|1-e^{-j\omega_1}\right| \\ &= (1-\cos\omega_1)^2+\sin^2\omega_1 \\ &= \sqrt{2}\sin\omega_1\end{aligned} \tag{5.33}$$

The output of the *contrast feature* itself, will be the square of the *filter* response (5.33) — due to the operator's *mean square* function. Thus the CON feature is a high pass filter and energy measure; the directionality of the former being controlled by displacement vector **d.**

Changing the size of **d** effectively changes the sampling frequency of the filter. For instance a vector **d** = (2,0) reduces the effective sampling frequency by half and changes the operator into a bandpass filter with weights (-1,0,1). Hence the frequency response of the CON(2,0) *filter* is :

$$\left|H_{CON}(\omega_1,\omega_2)\right| = 2\sin 2\omega_1 \tag{5.34}$$

Popular values for **d** are (1,0), (0,1), (1,1), (1,-1), (2,0), (0,2), (2,2) and (2,-2). Thus the contrast feature is in fact formed from a family of directional highpass and bandpass filters. This is confirmed by the theoretical and actual responses of the operator shown in figure 5.14.

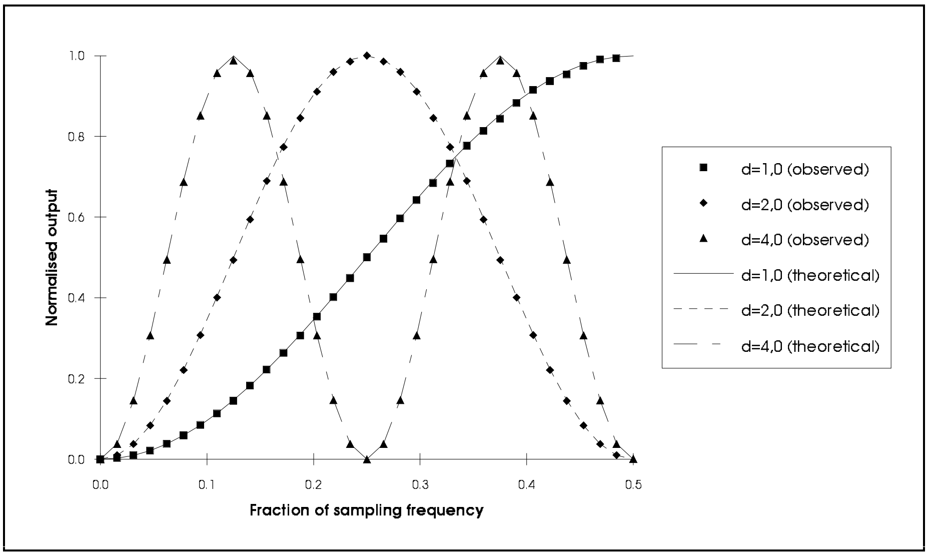

Figure 5.14 - CON operator : effect of changing the size of the displacement vector

b) Other co-occurrence features

The other three features ASM, ENT, and COR, have non-linearities which make their analysis in the frequency domain difficult. They were therefore only investigated empirically. The figure below shows one-dimensional plots of these three features.

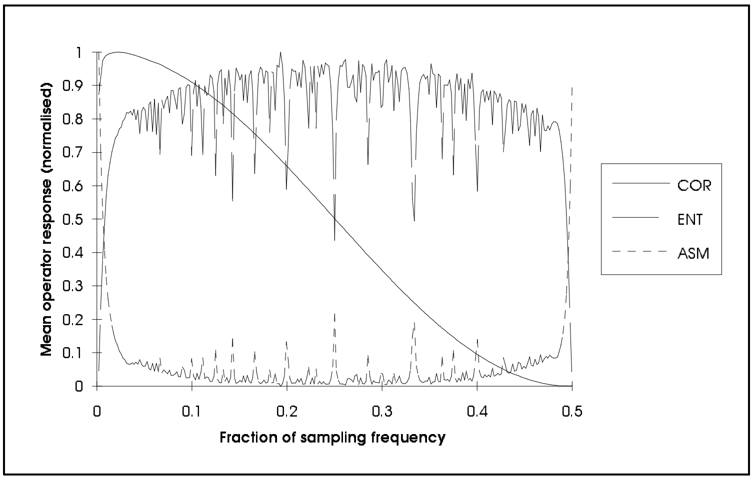

Figure 5.15 - One-dimensional frequency responses of co-occurrence features (d=(1,0), Ng = 16)

The above plots were obtained by running the co-occurrence features on strips of sine wave aligned parallel to the *x*-axis, and taking the average of the output. They show that the COR feature is a low pass filter. The frequency responses of the ENT and ASM operators are, in contrast, irregular. These "irregularities" are caused by sampling effects. When the period of the test sine wave is an integer multiple of the sampling period (which occurs at relative frequencies of 1/2, 1/3. 1/4 etc.) there are only a small number of unique grey-levels in the image. Thus at these frequencies the co-occurrence matrix is sparsely populated by a few high values of occurrences. Hence the ASM operator, which calculates the sum of squares of these occurrences (5.20), gives a high output value. The ENT operator is also a sum of a non-linear function of occurrences, and thus exhibits a similar behaviour. Since the frequency responses of these two operators are an extreme function of sampling effects, they will not be considered further here.

c) Two dimensional frequency responses

Co-occurrence features are essentially one-dimensional operators, in which the direction of the axis of the single dimension, is specified by the angle of the displacement vector **d.** Thus it is reasonable to expect the two-dimensional frequency response of an operator, to be the product of the operators' one-dimensional response θ_d and a unity gain element in the orthogonal direction. The figures below show this to be the case.

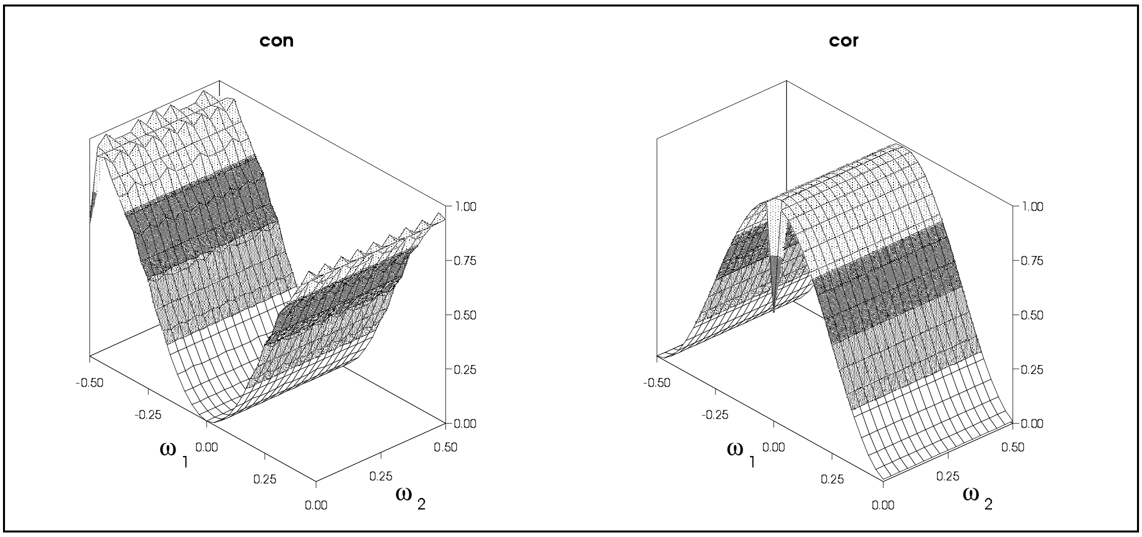

Figure 5.16 - Two-dimensional frequency responses of the co-occurrence contrast and correlation features for a displacement vector **d** = (0,1)

The frequency responses displayed above, were generated using images of corrugated sinusoids of the required x and y frequencies. They show that the CON and COR features are highly directional. They are therefore likely to be affected by illuminant tilt and slant, in a similar manner to Laws' uni-directional energy masks.

5.2.3. Tilt angle response

The illuminant tilt angle response was obtained by applying co-occurrence feature measures to the same test-set used in the Laws experiments. Unlike the Laws' operators they have no averaging filter (such as the ABSAVE macro statistic) rather the co-occurrence matrices are calculated directly from large windows. Indeed, such is the cost of these features, that they are most commonly calculated on large (e.g. 64x64) non-overlapping windows [Haralick73] or on the texture samples as a whole. Thus they are normally used either for classification of whole images or for very coarse segmentation. Few papers report their use for pixel level segmentation - an exception being [duBuf90] in which the use of a 7x7 moving window is described. Here a 33x33 moving window was used : in order to match the context employed by the Laws' features (a 5x5 mask plus a 29x29 macro statistic). In common with other researchers [duBuf90] the number of grey-levels Ng, and hence the size of the co-occurrence matrix used, did not noticeably affect the response of the features. Experiments with 16 and 256 grey-levels for instance gave similar results. The cheaper Ng = 16 option was therefore employed.

a) Isotropic and uni-directional textures

The tilt angle response of the four co-occurrence operators was obtained in an identical manner to that used for the Laws' masks. The responses of CON and COR to the four isotropic textures are shown below : similarly for ENT and ASM, except that the responses obtained using the directional texture *card1* are also shown.

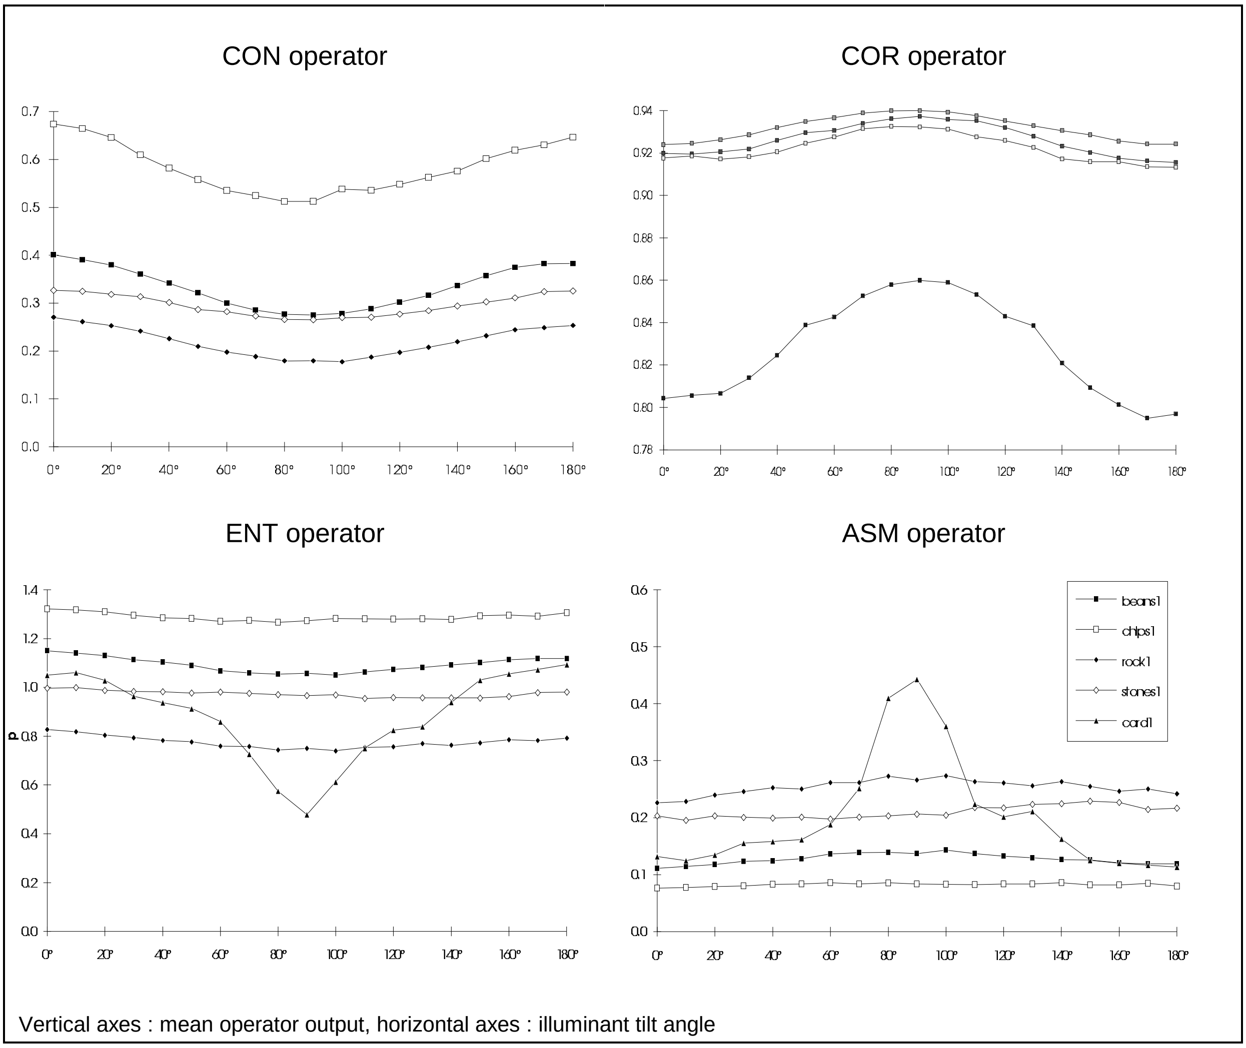

Figure 5.17 - Tilt angle response of co-occurrence features, Ng = 16, d = (1,0).

The above graphs show that the CON and COR features are sensitive to illuminant tilt. The former has a tilt angle response close to that of L5E5, which is not surprising given the similarity between their two-dimensional frequency responses (i.e. they both filter out frequencies in the direction $\theta = 90°$). However, the correlation feature COR produces almost the opposite angular response; having maxima at $\tau = 90°$ and minima at $\tau = 0°$, 180°. Since the same displacement vector, **d =** (1,0), was used for both operators, the results show that the direction of **d** cannot be used in isolation to predict the form of the tilt response. Examination of the operators' frequency responses shows why — CON(1,0) attenuates frequency components with an angle θ close to 90°, whereas COR(1,0) amplifies them.

In contrast the ENT and ASM features have relatively flat tilt responses when applied to the *isotropic* test textures (*beans1*, *chips1*, *rock1*, and *stones1*). However, when applied to the uni-directional texture *card1* — a corrugated surface in which the majority of frequency components run at $\theta = 0°$ — their tilt responses show that these operators are not invariant to tilt for all textures.

b) Normalisation

In section 5.1.2 the effect of image normalisation on the tilt response of Laws' operators was investigated, the motivation being that normalisation is used to compensate for lighting variations. Here therefore, its effect on the tilt response of co-occurrence features is described. Figure 5.18(a) below illustrates the COR feature's tilt response using images of the four isotropic textures, each normalised to a mean of 127 and a variance of 100.

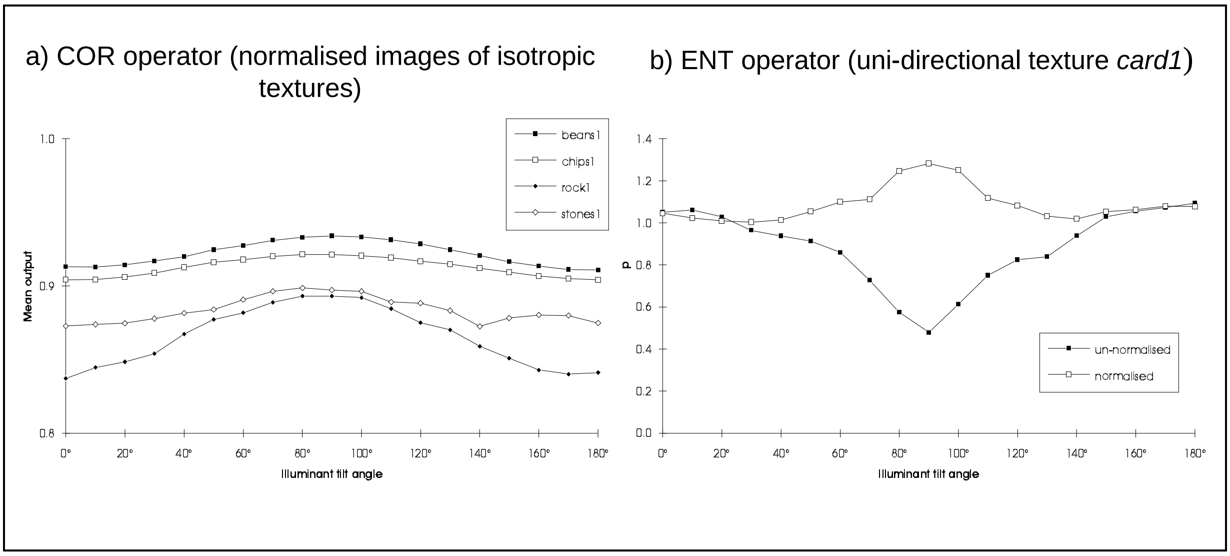

Figure 5.18 - The effect of normalisation on co-occurrence features

As was the case for Laws' operators, normalisation clearly does not compensate for variation in illuminant tilt when the surface textures are isotropic. However, for uni-directional textures a different response would be expected. For an isotropic texture variation in illuminant tilt does not affect the total variance of the image. Enhancement of components coincident with τ is compensated for by attenuation of components at 90° to τ. Hence normalisation will have the same effect on each image regardless of tilt. The variance of a uni-directional texture however, will be affected by changes in tilt; as the energy is concentrated in one direction. Thus normalisation may compensate for variation due to tilt when the texture is uni-directional. Figure 5.18(b) above shows the ENT tilt responses to normalised and un-normalised images of a uni-directional texture *card1*. In this case normalisation has significantly affected the tilt response, indeed for τ = 80° to 100° it has "overcompensated". This effect is addressed in section 5.3.2(b). What is clear however, is that unlike the isotropic case, normalisation of a uni-directional texture does significantly affect the tilt response; and that for the example *card1*, it has flattened it.

5.2.4. Slant angle response

Section 5.1.3 showed that normalisation of texture images, can help compensate for the effects of slant angle variation for Laws' L5E5 operator. The experiment was therefore repeated, to determine whether or not co-occurrence features may be similarly compensated. Figure 5.19(a) below shows that the ASM feature is not invariant to σ when used with un-normalised images.

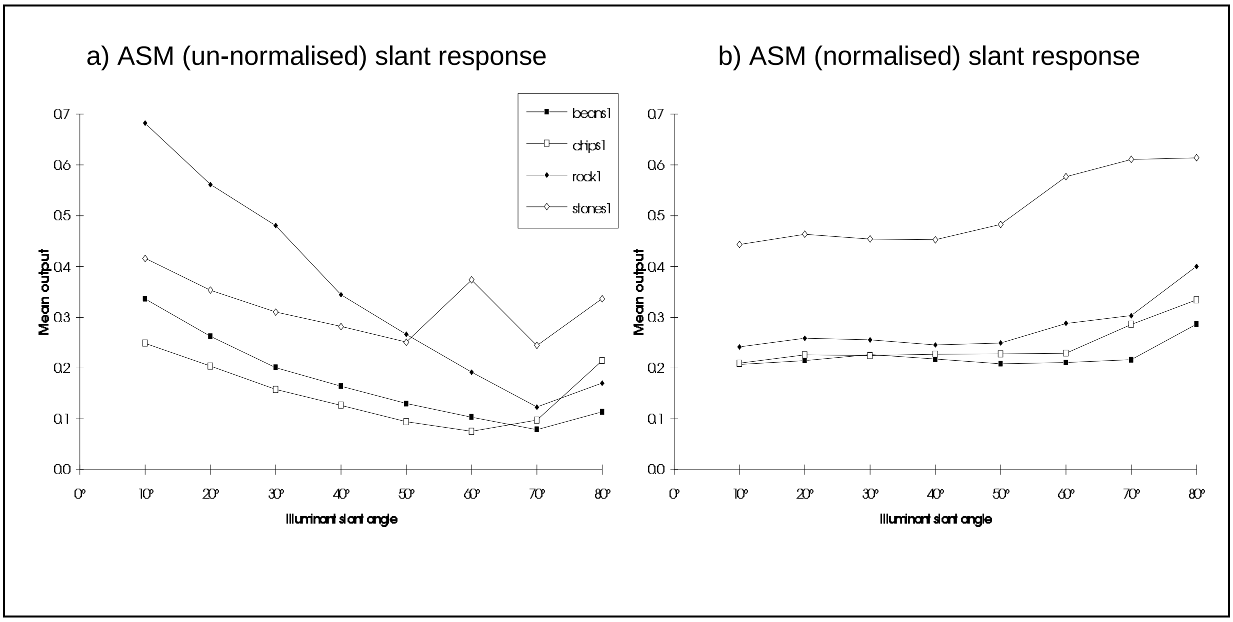

Figure 5.19 - The effect of image normalisation on the ASM *slant response.*

The second graph, 5.19(b), shows that normalisation has partly compensated for variation in illuminant slant. It has not been as successful as was the case for L5E5 (figure 5.13), as there are marked variations in the normalised response for σ greater than 60°. Nevertheless, normalisation has significantly reduced the variation in the operator's response for values between 10° and 50°. Note however, that this has again been purchased at the cost of reduced separations between test texture means.

5.2.5. Summary

This section has investigated the effect of variation in illuminant tilt and slant on co-occurrence features. To summarise :

- A formulation of the COR operator has been developed, which provides an efficient implementation in which co-occurrence matrices do not need to be maintained.
- The features' frequency responses show that the CON and COR operators are similar to Laws' masks, in that they are directional low, high, and band pass filters.
- These directional operators were shown not to be invariant to illuminant tilt.
- The ENT and ASM features were not invariant to tilt when applied to the *uni-directional* texture *card1*.
- Normalisation of images was shown to be able to compensate (in fact over-compensate) for tilt variation effects when applied to the uni-directional test texture, but it had little effect on isotropic test textures.
- Normalisation was shown to be capable of compensating for illuminant slant variations at lower angles (less that 60°), but it was also observed that it reduced the separations between the test texture means.

5.3. Linnett’s operator

Fractal dimension [Mandelbrot83] is an appealing concept to use as a basis for a texture feature as it has been suggested that it provides a measure of roughness [Pentland84] [Arduini92] [Dennis89]. A number of researchers have used estimates of fractal dimension for texture classification with mixed results [Pentland84] [Medioni84] [Keller89] [Mosquera92] [Peli90]. The main problems being the computational cost of and the limited classification accuracy available from this single feature measure. Linnett's operator [Linnett91a] does not in fact estimate fractal dimension — rather it utilises the information available from the first one, two, or three iterations, of an iterative process that does. Thus the computational complexity is greatly reduced. In addition, instead of using a single isotropic measure, Linnett used the masks shown below to produce seven operators each having a different directional characteristic.

X	X	X

m1

		X
	X	
X		

m2

	X	
	X	
	X	

m3

X		
	X	
		X

m4

	X	
X	X	X
	X	

m5

X		X
	X	
X		X

m6

X	X	X
X	X	X
X	X	X

m7

Figure 5.20 - The seven directional masks for Linnett's operator

Linnett's operator is based upon Peleg's iterative blanket method for estimating the fractal dimension of a surface [Peleg84]. Peleg's algorithm creates a series of upper and lower surfaces each of which is a radius of λ from the previous upper or lower surface, respectively.

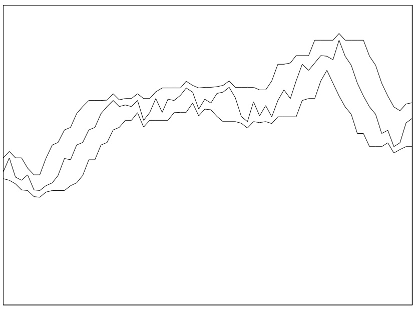

Figure 5.21 - Sections from example surfaces of the blanket method : original surface (middle), upper blanket (top), and lower blanket (bottom)

It is the scaling behaviour of the volume enclosed between upper and lower pairs that yields the estimate of fractal dimension. Linnett however, uses the enclosed volume directly as a texture feature and therefore avoids the requirement for repeated iterations. Thus his operator does not provide an estimate of fractal dimension.

For iteration n the upper $u_n(x, y)$ and lower $l_n(x, y)$ blankets are defined as :

$$u_n(x, y) = \max\{u_{n-1}(x, y) + \lambda, u_{n-1}(x + i, y + j)\}\forall(i, j) \in m \tag{5.35}$$

$$l_n(x,y) = \min\{l_{n-1}(x,y) - \lambda, l_{n-1}(x+i, y+j)\} \forall (i,j) \in m \qquad (5.36)$$

where

m is a neighbourhood of pixels defined by one of the masks shown in figure 5.20,

$u_1(x,y) = l_1(x,y) = I(x,y)$ = the original image.

The feature images themselves are derived from the volume enclosed locally by the upper and lower blankets :

$$v_n(x,y) = u_n(x,y) - l_n(x,y) \qquad (5.37)$$

Thus Linnett's operator is similar to Dinstein's *maxdif* (maximum difference) operator [Dinstein84] — its key feature being its directional characteristics as defined by the seven masks.

5.3.1. Frequency response

a) The one dimensional case

Linnett examined the one dimensional frequency response of his operator both theoretically and empirically for a single iteration [Linnett91a]. He determined that it would behave as any one of six non-recursive filters depending upon the data. The frequency responses of these six modes of operation are :

$$v_1(x) = I(x-1) - I(x) \Rightarrow H(\omega) = -1 + \cos\omega - i\sin\omega \qquad (5.38)$$

$$v_2(x) = I(x-1) - I(x+1) \Rightarrow H(\omega) = -i2\sin\omega \qquad (5.39)$$

$$v_3(x) = I(x) - I(x-1) \Rightarrow H(\omega) = 1 - \cos\omega + i\sin\omega \qquad (5.40)$$

$$v_4(x) = I(x) - I(x+1) \Rightarrow H(\omega) = 1 - \cos\omega - i\sin\omega \qquad (5.41)$$

$$v_5(x) = I(x+1) - I(x-1) \Rightarrow H(\omega) = i2\sin\omega \qquad (5.42)$$

$$v_6(x) = I(x+1) - I(x) \Rightarrow H(\omega) = -1 + \cos\omega + i\sin\omega \qquad (5.43)$$

He observed from the above that there are in fact only two different magnitude responses:

$$|H_1(\omega)| = 2\sin\omega \qquad (5.44)$$

$$|H_2(\omega)| = \sqrt{2 - 2\cos\omega}, \qquad (5.45)$$

Linnett verified these theoretical results by measuring the response of each of the six modes (5.38) to (5.43) to a sine wave with frequencies up to the Nyquist limit. Note that the operator was not used in the normal way but *fixed* into the operating mode under investigation. What is of greater interest here though is the operator's frequency response per se. Hence the graph below shows the observed frequency response of the operator itself without any restrictions as to its modes of operation. For comparison the magnitude frequency responses of (5.44) and (5.45) above have also been plotted on this graph.

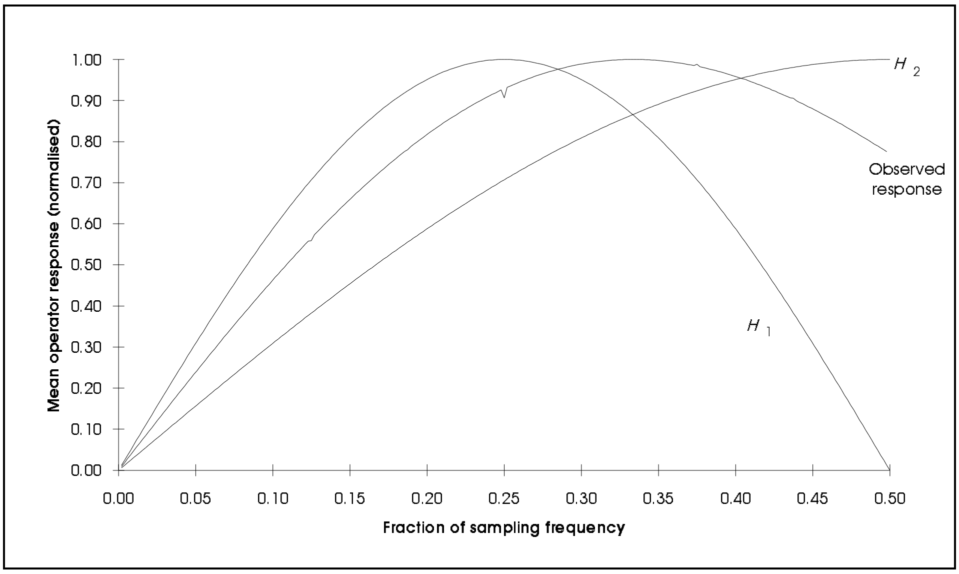

Figure 5.22 - Observed one dimensional frequency response of Linnett's operator, together with the two theoretical cases

The figure above shows that the observed frequency response is a combination of the two theoretically derived cases. It is a type of band-pass filter and would therefore be expected to have a similar performance to Laws' S5 mask or the co-occurrence CON texture measure.

b) The two-dimensional case

The directional characteristics of the operator are controlled by the mask that it is used with. Linnett used seven masks shown above in figure 5.20. Masks *m1* to *m4* are uni-directional and their two-dimensional response would be expected to be a straight projection of the one-dimensional case. On the other hand Masks *m5* to *m7* are multi-directional and their responses are more difficult to predict. The figure below shows the responses of six of the seven versions of the operator to corrugated sinusoidal images with spatial frequencies up to the Nyquist limit. The frequency response of *m4* has not been shown, it is simply a 90° rotation of *m2*.

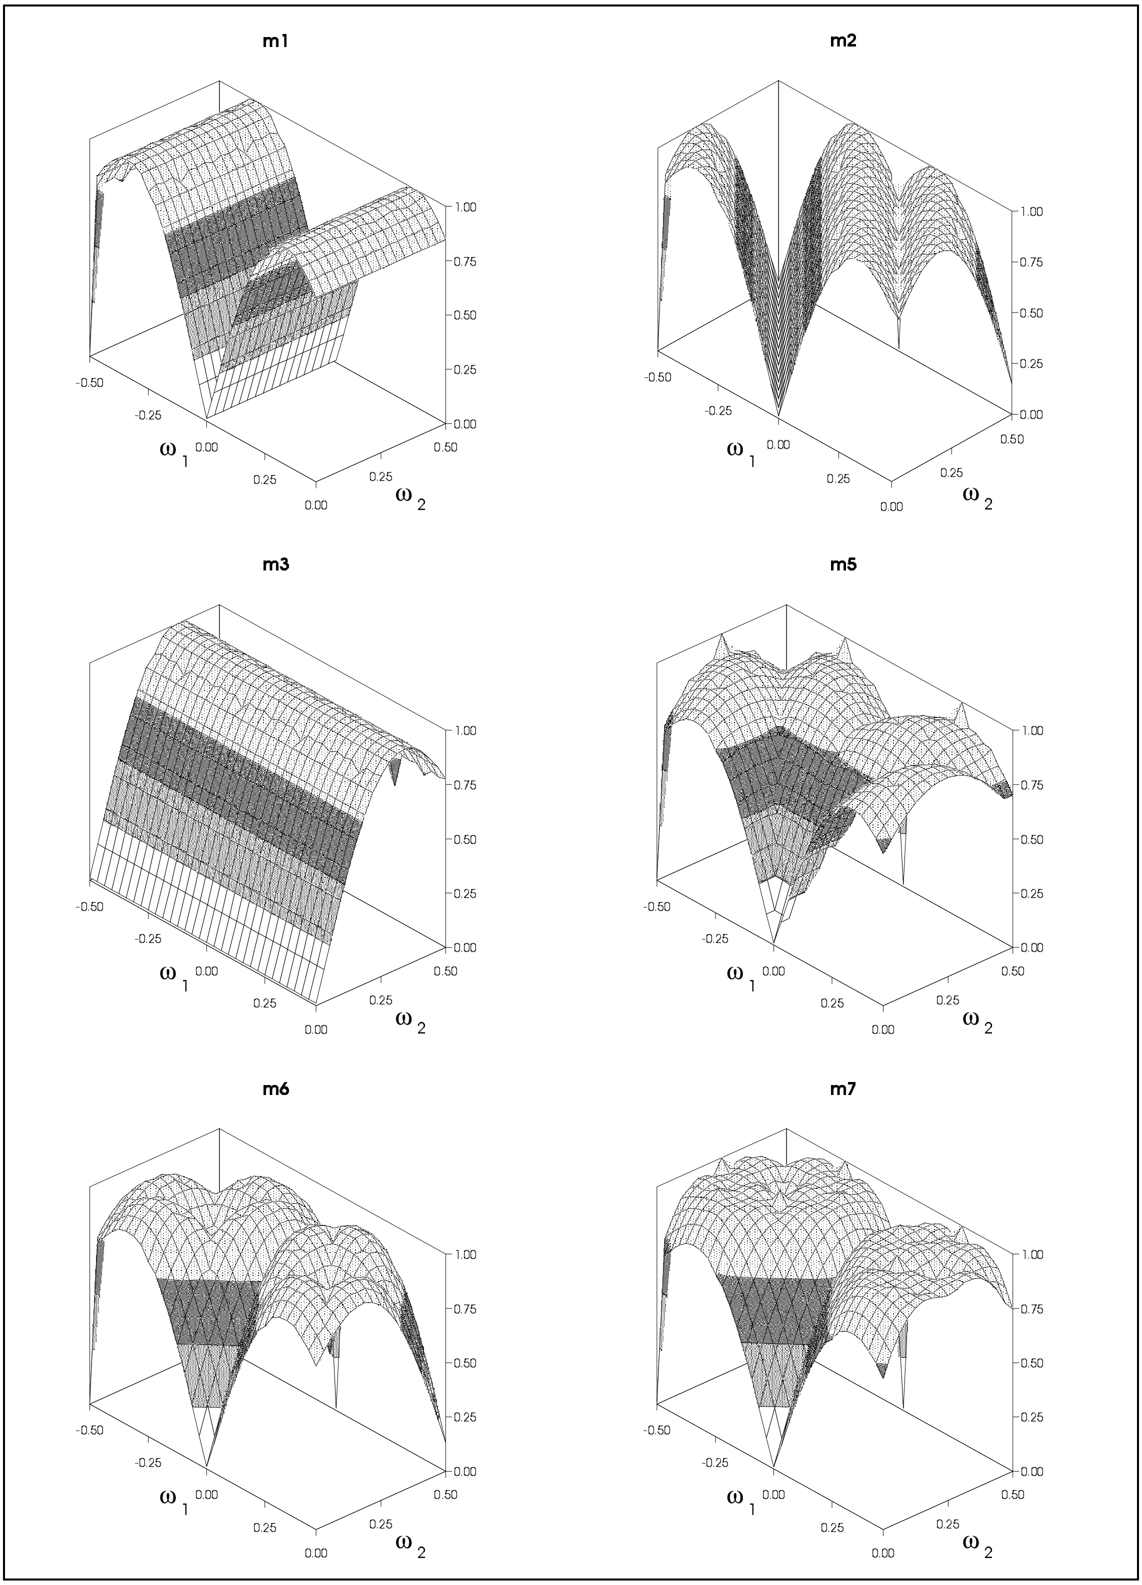

Figure 5.23 - Linnett's operator : observed two dimensional frequency responses

The above show that *m1*, *m3*, *m2* (and hence *m4*) are highly directional, whereas *m5*, *m6*, and *m7* are largely isotropic. The first four masks are therefore likely to be affected by illuminant tilt in a similar manner to Laws' L5E5 operator, i.e. they are not anticipated to be invariant to tilt even when used on isotropic textures. The other three masks would be expected to be tilt invariant for isotropic texture but not for directional texture.

5.3.2. Tilt angle response

Tilt angle responses of the seven masks were obtained using one directional and four isotropic textures. A single iteration of each version of the operator was used. The experimental procedure was as used to obtain the tilt response of Laws' filters, the only difference being that, as the seven masks are based on 3x3 kernels, a 31x31 ABSAVE macro statistic was used to keep the overall size of the local neighbourhood at 33x33.

a) Un-normalised tilt response

Figure 5.24 depicts the responses of four of the masks to un-normalised images of the four isotropic textures (*beans1*, *chips1*, *rock1*, *stones1*) and the directional texture (*card1*).

Directional masks

Mask *m1* is clearly not invariant to variation in illuminant tilt and has a very similar response to Laws' L5E5 operator, which is explained by the similarity between their low frequency responses. (Note that the lower frequencies are likely to dominate the above responses due to the exponential nature of the textures' PSDs.) With respect to the isotropic textures; masks *m2*, *m3*, and *m4*, give similar results to *m1* — rotated by the appropriate angle. For example the maxima in *m2*'s response occur at 45°. Note however that *m3*'s response to *card1* is in contrast with the other three directional masks almost flat — this is due to it being insensitive to spatial frequencies at 0°, and hence it detects very little energy from the vertical corrugations.

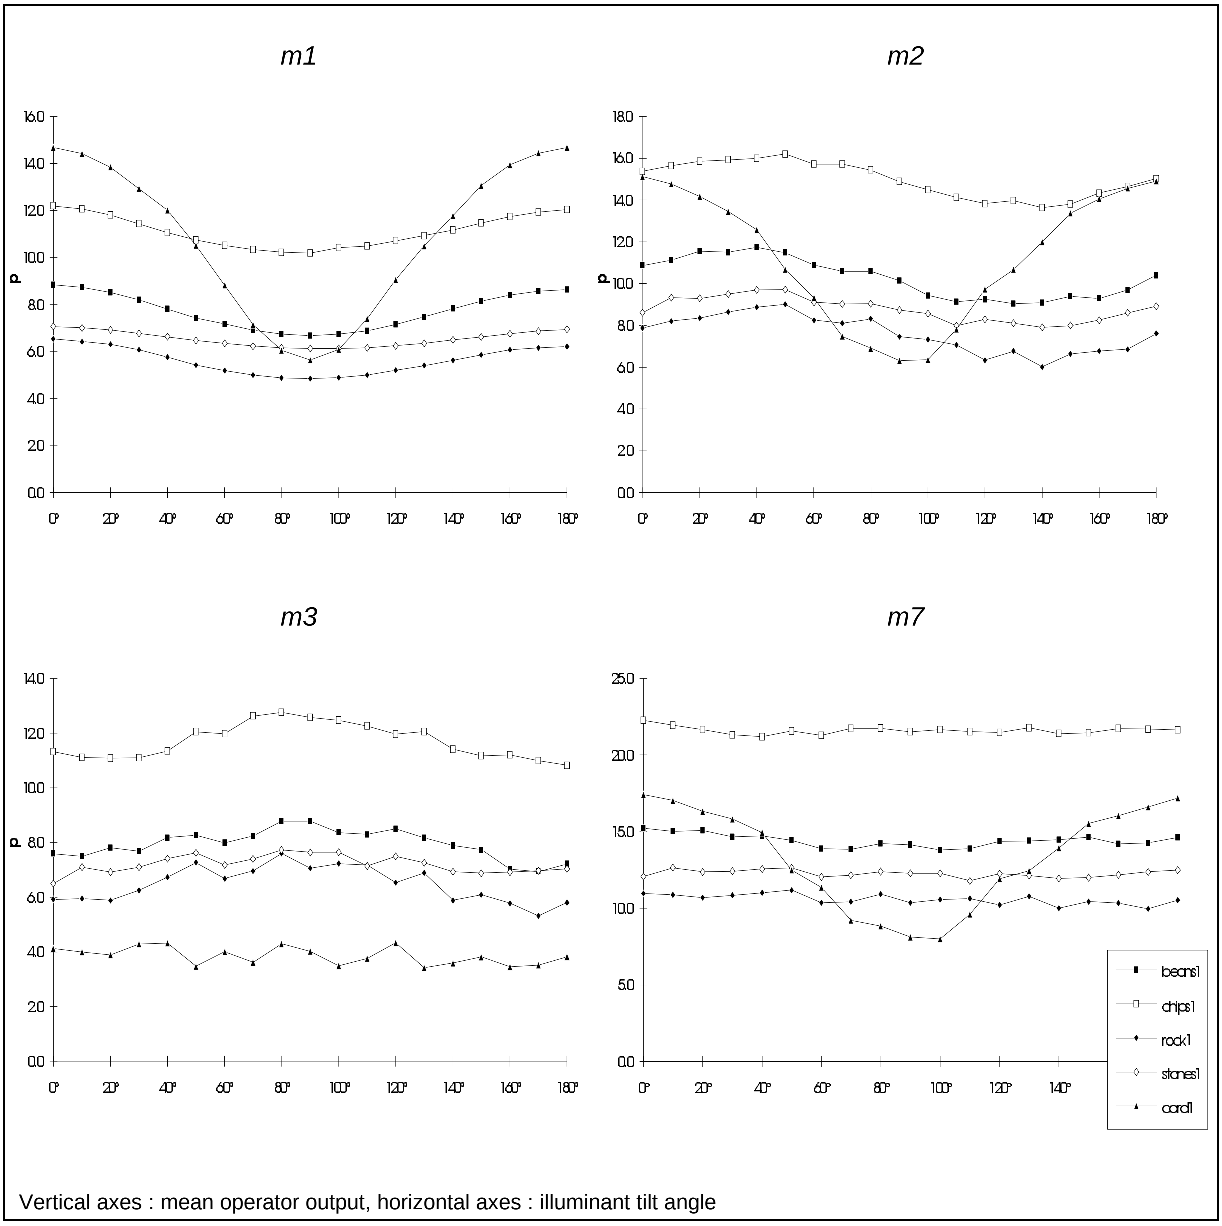

Figure 5.24 - Linnett's operator (un-normalised) : tilt response

Isotropic masks

The frequency responses of *m5*, *m6*, and *m7* have shown them to be largely isotropic. It is not surprising therefore that their tilt responses are (i) very similar to each other, and (ii) invariant to changes in tilt when applied to *isotropic* texture. Consequently only *m7*'s tilt response is shown above in figure 5.24.

The flat responses of isotropic operators to isotropic textures occur because directional illumination enhances texture components coincident with the illuminant tilt but attenuates those components at right angles to it. Unfortunately the same is not true of directional textures. Uni-directional textures will be attenuated by the imaging process when the illuminant tilt is perpendicular to the texture's direction, and there will be no compensation from amplified frequency components co-incident with the illuminant's direction of tilt because there are none present in the texture. An example of this effect is *m7*'s tilt response to the uni-directional texture *card1* - a corrugated surface in which the majority of frequency components run at $\theta = 0°$. This response shows that for uni-directional textures the operator is certainly not invariant to tilt. Masks *m5* and *m6* gave similar responses to this texture.

b) Normalised tilt response

The illuminant tilt angle response of Linnett's operator using normalised images was investigated as before. The aim being to determine whether or not such pre-processing compensates for variation in τ. Figure 5.25 below shows the response of the operator using *m1*.

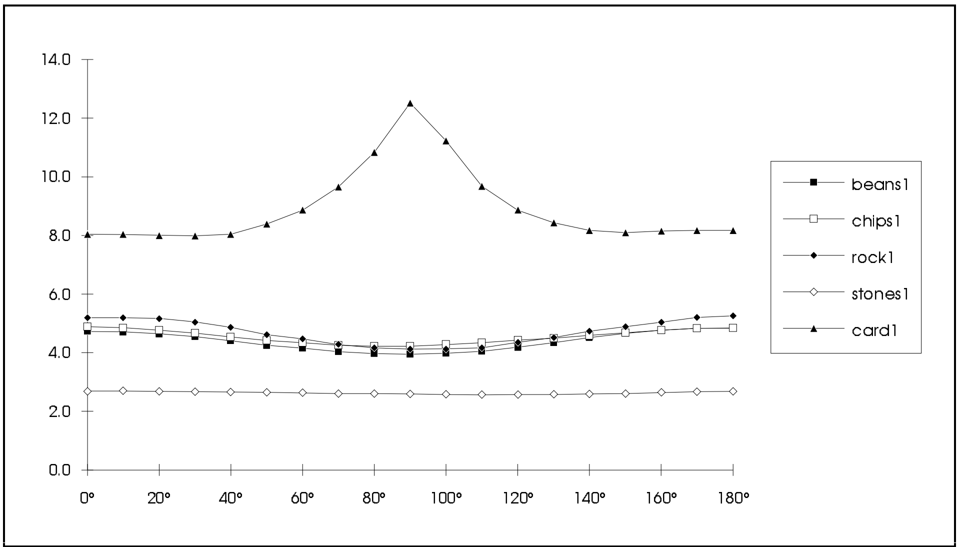

Figure 5.25 - The effect of normalisation on mask m1's tilt response

The shape of the responses of the four isotropic textures have slightly flattened but are still clearly affected by tilt. The response of the directional texture *card1* has in contrast been markedly affected. Indeed normalisation appears to have "over compensated" for angles of τ around 90°. An explanation for this behaviour is as follows.

From the Lambertian image model (2.1) :

$$I(x,y)=\frac{-p\cos\tau\sin\sigma\ -\ q\sin\tau\sin\sigma\ +\ \cos\sigma}{\sqrt{p^2\ +\ q^2\ +\ 1}}$$

$$=\left(-p\cos\tau\sin\sigma\ -\ q\sin\tau\sin\sigma\ +\ \cos\sigma\right)\left(1-\frac{\left(p^2+q^2\right)}{2!}+\frac{9\left(p^2+q^2\right)^2}{4!}\cdots\right) \tag{5.46}$$

Now, for a directional texture that contains components only in the direction θ = 0°

$$q=\frac{\partial N_H}{\partial y}=0 \tag{5.47}$$

Hence, for an illuminant tilt angle of τ = 0°

$$I_{\tau=0^\circ}(x,y)=\left(-p\sin\sigma+\cos\sigma\right)\left(1-\frac{p^2}{2!}+\frac{9p^4}{4!}\cdots\right) \tag{5.48}$$

and for τ = 90°

$$I_{\tau=90^\circ}(x,y)=\cos\sigma\left(1-\frac{p^2}{2!}+\frac{9p^4}{4!}\cdots\right) \tag{5.49}$$

By comparing (5.48) and (5.49) it can be seen that changing the illuminant tilt from 0° to 90° removes the proportional and other odd terms ($-p\sin\sigma$ etc.). Thus compared with its τ = 0° counterpart, the τ = 90° image will not contain the fundamental frequency or odd harmonics. Now if both images are normalised to have the same variance, the net result of the change of tilt angle from 0° to 90°, will be a shift of power from the odd to the even harmonics. Figures 5.26 and 5,27 show that this is indeed what happens.

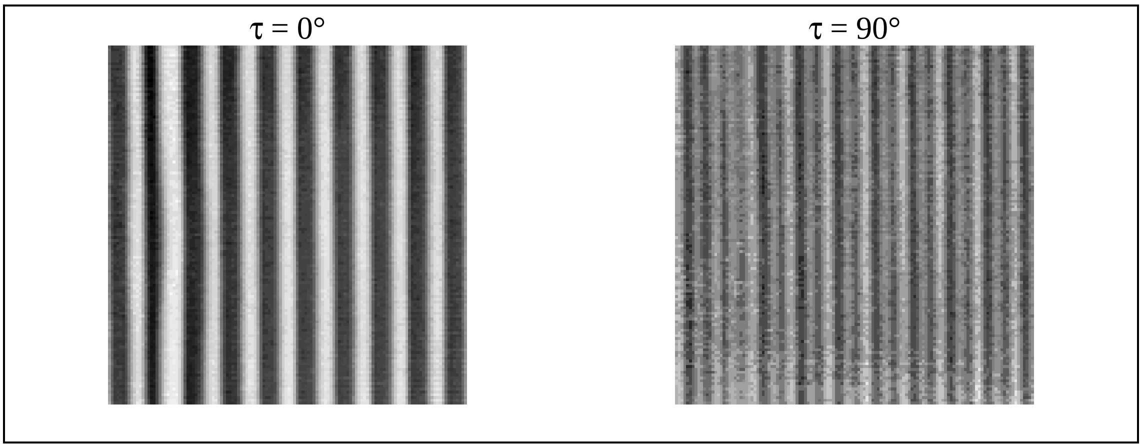

Figure 5.26 - Normalised 128x128 samples of card1 at illuminant tilt angles of 0° and 90° (frequency of corrugations ≅ 0.08 times the sampling frequency)

From figure 5.27 it can be seen that (i) the *m1* operator is more sensitive to the 2nd harmonic ($\omega \cong 0.16\omega_s$) than the fundamental ($\omega \cong 0.08\omega_s$), and that a tilt of 90° reduces the first and amplifies the second.

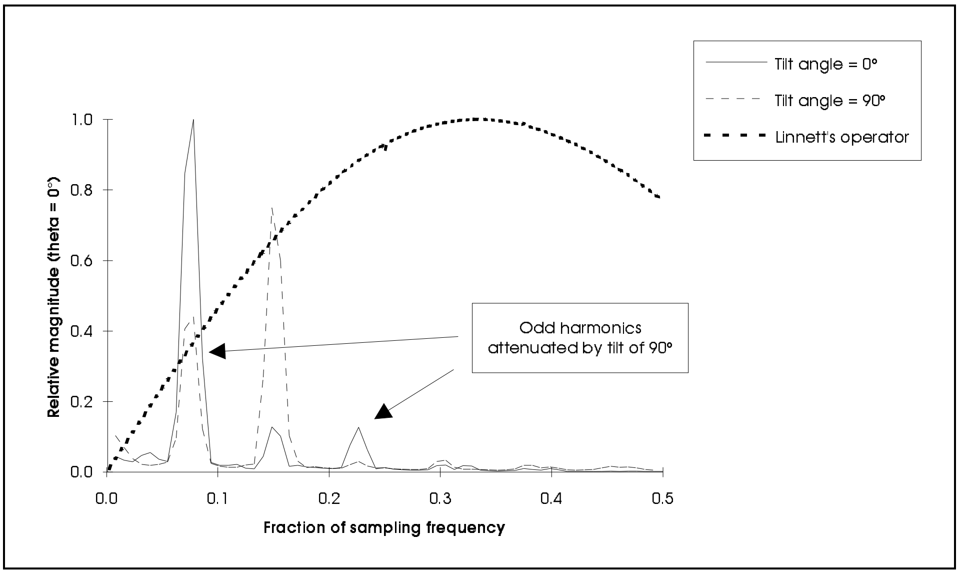

Figure 5.27 - Radial sections of magnitude spectra (at θ = 0°) of card1 and operator m1

Thus the end result is that the frequencies to which Linnett's *m1* operator are more sensitive are boosted in normalised images as the illuminant tilt approaches an angle of 90° : hence the "over compensation" effect.

An "over-compensation" effect of normalisation is also apparent, although less obvious in the slant response described in the next section.

5.3.3. Slant angle response

Previous sections of this chapter have shown that normalisation of images can help compensate for the effects of variation in illuminant slant (σ) on Laws' and co-occurrence texture features. This section therefore addresses the same issue for Linnett's operator. The slant responses were obtained using the same test set of normalised and un-normalised image textures as used in the Laws' and co-occurrence experiments. Mask *m6*'s responses are depicted below.

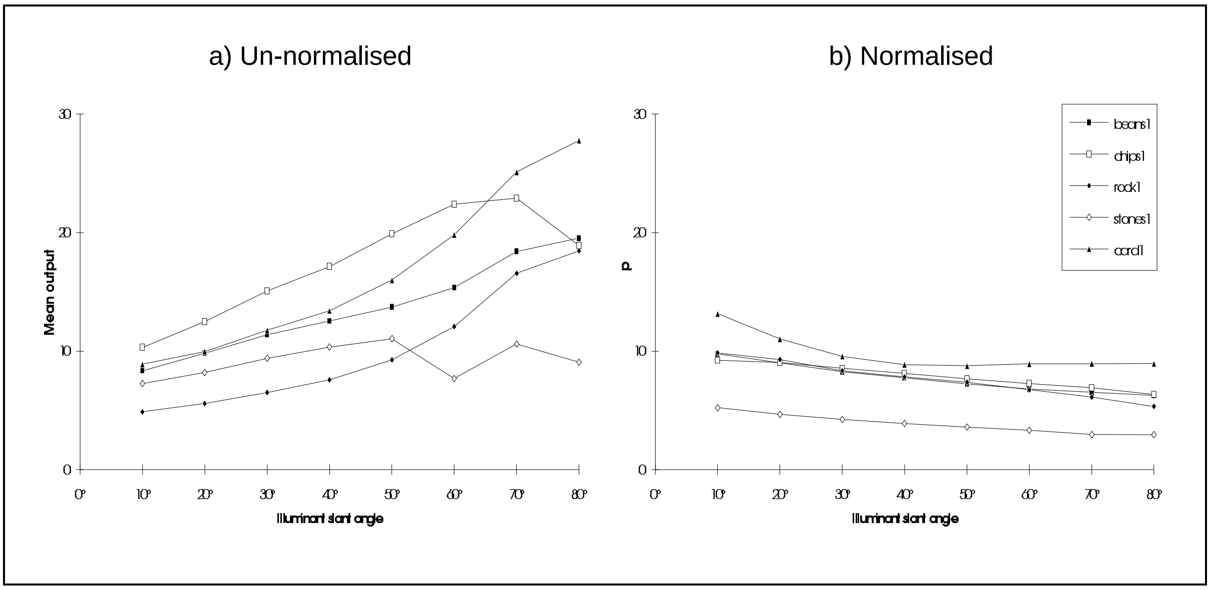

Figure 5.28 - The effect of normalisation on mask m6

It is apparent from the above graphs that normalisation has helped reduce the variation of mask 6's output, but that it has again "over-compensated": the mean output at low angles of slant σ being greater than that of higher angles. As the slant angle approaches the vertical, that is $\sigma \to 0$, the proportional and other odd terms in (5.46) again tend to zero, giving a similar effect to the "over compensation" of tilt variation discussed previously. Note that the separation between class means has again been reduced.

5.3.4. Summary - Linnett's operator

This section has investigated the effect of variation in illuminant tilt and slant on Linnett's operator. The main points to emerge from this section are as follows.

- The directional operators *m1* to *m4* are not invariant to variation in tilt, with respect to either the normalised or un-normalised test textures.
- Masks *m5, m6,* and *m7,* which have approximately isotropic frequency responses, are not affected by tilt when applied to the isotropic test textures, but are affected when applied to the directional texture *card1*.
- Normalisation of images of the directional texture *card1* over-compensates for variations in tilt at tilts of around 90° to the texture direction.
- Normalisation of images of both isotropic and directional test textures does compensate for slant angle variation (although some over compensation is evident).

However, this effect is provided at the expense of reduced separation between feature means.

5.4. Metrics for class separation and sensitivity to illuminant variation

The three preceding sections of this chapter have discussed qualitatively the effects of illuminant variation. This is sufficient to explain the general phenomena observed. The development of a *quantitative* measure of these effects would however, facilitate feature selection and provide a valuable tool to extend the previous discussion. The Mahalanobis' distance [Tou74] is commonly used to provide a measure of the ability of a feature set to separate two classes. It uses separation between the class means adjusted by a factor to account for classes' variances. For a single feature this measure reduces to the "generalised difference" [Davis73] :

$$D^2 = \frac{(\mu_1 - \mu_2)^2}{\sigma_p{}^2} \tag{5.50}$$

where

$\sigma_p{}^2$ is the pooled variance [Davis73] of the two classes concerned, and

μ_1, μ_2 are the means of the feature measure's outputs for the two classes.

In cases where more than two classes are involved it is normal to compute D^2 for each possible class pair and choose the worst case (lowest) result.

Since the Mahalanobis distance provides a measure of the separation between two distributions, it may be adapted to provide an illuminant sensitivity metric. That is, it may be used to measure the maximum displacement of an operator's output distribution caused by a change in illumination. Thus a measure of the tilt sensitivity of a feature, with reference to a given texture, may be computed as follows :

(i) Obtain feature images of the texture over the required range of illuminant tilt.

(ii) Determine D^2 for each possible pair of feature images.

(iii) Choose the worst case (highest) result.

The above is an expensive procedure : a cheaper alternative is to calculate D^2 between the feature images having the highest and lowest mean values. This will yield similar results to the above if image feature variance does not change significantly with illuminant tilt. Thus the tilt sensitivity metric used here is defined as

$$D_\tau^2 = \frac{2(\mu_{\max} + \mu_{\min})^2}{\sigma_{\max}^2 + \sigma_{\min}^2} \tag{5.51}$$

where

> $\mu_{\max}$ and $\mu_{\min}$ are the maximum and minimum mean operator outputs over the required range of illuminant tilt, and
>
> $\sigma_{\max}^2$ and $\sigma_{\min}^2$ are the variances of the operator's output at tilt angles at which $\mu_{\max}$ and $\mu_{\min}$ occur.

For a number of textures the mean tilt sensitivity across texture types provides a single figure of merit. Thus the *mean tilt sensitivity* referred to in table 5.1 is defined as

$$\overline{D_\tau^2} = \frac{1}{n} \sum_{\substack{\text{For each} \\ \text{texture}}} D_\tau^2 \tag{5.52}$$

where

> *n* is the number of textures

Table 5.1 contains the results of applying this metric to co-occurrence, Laws', and Linnett's features. Two texture sets were used : *set a* comprising the four isotropic test textures, and *set b* consisting of *set a* together with the directional texture *card1*. Both the original 512x512 floating point images (512f) and their normalised versions (512N) were used. For convenience the first five rows of the table contain means of five different groups of features, collected by feature type (Laws, Linnett, or co-occurrence) and normalisation.

			Mean tilt sensitivity $\left(\overline{D_\tau^{\ 2}}\right)$		Class separation (D^2)	
			set a	*set b*	*set a*	*set b*
Mean co-occurrence			0.80	9.33	0.09	0.05
Mean Laws			1.14	9.17	1.01	0.37
Mean Linnett			1.46	8.96	0.15	0.11
Mean 512f			1.22	10.67	0.58	0.24
Mean 512N			1.12	7.99	0.64	0.24
Cooccurrence	ASM	512f	0.32	3.43	0.19	0.11
Cooccurrence	CON	512f	2.02	8.22	0.12	0.09
Cooccurrence	ENT	512f	0.47	11.79	0.32	0.12
Cooccurrence	COR	512f	0.76	32.72	0.01	0.01
Cooccurrence	ASM	512N	0.28	1.28	0.02	0.01
Cooccurrence	ENT	512N	0.26	1.80	0.05	0.03
Cooccurrence	CON	512N	1.51	6.43	0.01	0.01
Cooccurrence	COR	512N	0.77	9.00	0.04	0.02
Linnett	*m1*	512f	2.20	17.84	0.19	0.14
Linnett	*m2*	512f	2.81	12.94	0.17	0.10
Linnett	*m3*	512f	2.71	2.46	0.11	0.11
Linnett	*m4*	512f	2.00	12.46	0.23	0.14
Linnett	*m5*	512f	0.43	10.74	0.26	0.19
Linnett	*m6*	512f	0.25	9.95	0.34	0.23
Linnett	*m7*	512f	0.32	9.84	0.28	0.20
Linnett	*m1*	512N	1.53	8.76	0.03	0.03
Linnett	*m2*	512N	2.74	7.61	0.03	0.03
Linnett	*m3*	512N	2.70	5.03	0.09	0.05
Linnett	*m4*	512N	1.89	6.91	0.04	0.04
Linnett	*m5*	512N	0.40	5.86	0.12	0.12
Linnett	*m6*	512N	0.22	8.10	0.04	0.04
Linnett	*m7*	512N	0.27	6.99	0.09	0.09
Laws	*E5E5*	512f	0.16	0.15	0.00	0.00
Laws	*E5S5*	512f	0.41	0.35	0.02	0.02
Laws	*L5E5*	512f	2.55	18.03	0.28	0.21
Laws	*L5S5*	512f	2.90	12.94	0.15	0.14
Laws	*R5R5*	512f	0.05	0.10	0.01	0.00
Laws	*E5E5*	512N	0.19	9.40	0.06	0.05
Laws	*E5S5*	512N	0.30	9.82	1.26	0.41
Laws	*L5E5*	512N	1.95	2.56	0.08	0.08
Laws	*L5S5*	512N	2.62	26.74	0.18	0.18
Laws	*R5R5*	512N	0.27	11.62	8.03	2.58

Table 5.1 - Tilt sensitivity and class separation of Laws', Linnett's, and co-occurrence features.

Table 5.1 shows that :

(i) Normalisation significantly reduces the tilt sensitivity of features derived from *set b* which contains the directional texture *card1*.

(ii) Normalisation does not significantly affect the tilt sensitivity of *set a* which contains only isotropic textures.

(iii) Features with frequency responses which are approximately omnidirectional or bi-directional have low tilt sensitivities when used with isotropic textures (*set a*). See for instance Linnett's *m4, m5, & m7,* Laws' R5R5 & E5E5, and co-occurrence features ENT and ASM.

(iv) Conversely features with uni-directional frequency responses have high tilt sensitivities and thus Laws L5S5, E5S5, Linnett's *m1* to *m4*, and the co-occurrence feature CON, would all provide discrimination between differing illumination tilts. Thus they would be useful in tilt estimation schemes.

(v) For the texture test set employed Laws' energy masks provide on average the best potential class separation. The next best is provided by Linnett's operator.

5.5. Conclusions

This chapter has examined the effects of variation of illuminant slant (σ) and tilt (τ) on three sets of texture features. Of particular interest was the effect of normalisation; as it had been suggested in chapter 2 that normalisation of images could compensate for slant variation but not for tilt variation (except where uni-directional images are concerned). The behaviour of each feature set was therefore investigated. Images of four isotropic textures and one uni-directional texture were captured under a range of illuminant slant and tilt conditions. These data sets were presented to the feature measures and the effects on the resulting output distributions, in terms of means and histograms, were recorded. In addition a new tilt sensitivity metric, based upon the Mahalanobis distance, was developed and used to assess the effect of tilt variation on these features.

To summarise, the main conclusions drawn from the preceding investigations with respect to the limited set of test textures employed, are as follows.

- The effect of change in illuminant tilt angle was shown to alter according to the directional characteristics of the test texture and the feature measure concerned. The following table summarises these findings :

	Isotropic features	**Bi-directional features**	**Uni-directional features**
Isotropic texture *(beans1, chips1, rock1 & stones1)*	Not significant	Affected	Significantly affected
Uni-directional texture *(card1)*	Significantly affected	Significantly affected	Significantly affected

Table 5.2 - The effect of illuminant tilt on directional and isotropic feature measures

- Normalisation was shown to reduce the tilt sensitivity of features when applied to the directional texture *card1*.
- It was also shown that normalisation does not significantly affect the tilt sensitivity of features when applied to the isotropic test textures.
- Variation in illuminant slant has been shown to significantly affect each of the three feature sets.
- Image normalisation has been shown to reduce this variation (at a cost of an associated reduction in separation between class means).

Chapter 6

Classification

The previous chapter has shown that variation of the illumination's direction can affect the outputs of Laws', Linnett's' and co-occurrence feature sets. Such variations may be encountered in a variety of situations (as described in chapter 1). Unfortunately the effect of illuminant variation on classification accuracy is likely to be dependent on the application. This is because alterations in lighting effect a movement of class members in feature space. But these displacements are only significant if they cross decision surfaces, the position of which are dependent upon the characteristics of the original training set. Hence the number of classification errors caused by a change in lighting is a function of the feature set selected, the number of textures, the characteristics of the textures, and of course the illuminant variation itself. Thus the effects of illuminant variation can only be assessed with respect to a particular classification task.

It is not within the scope of this chapter to identify and test classification tasks representative of all of the applications mentioned in chapter 1. Rather the aims of the following sections are :

(i) to show that lighting variation can significantly affect *a* classification task,

(ii) to develop a prototype compensation scheme, and

(iii) to show that such a compensation scheme can reduce the effects of illuminant variation for the chosen classification task.

Thus the purpose of the work described in this chapter was not to develop an optimum classifier, but rather it is to investigate the effects of illumination variation on a *representative* classifier. Effort was *not* invested in feature selection — the feature sets used were selected purely on the basis of popularity in the literature and ease of implementation. Nor was any post-processing, such as mode filtering [Greenhill93], employed.

The investigation reported here, uses one model classification example throughout. The job consists of classifying montages of four textures under varying illuminant tilt and slant angles. That is either the tilt or slant angle is varied *between* training and classification sessions. Thus the first objective of this chapter is to assess the effect of illuminant variation on this model task. The second objective is to investigate the effect of image normalisation — as previous chapters have suggested that normalisation may reduce classification errors that are due to variations in illuminant direction. The third and final objective is the development of a prototype tilt-compensation scheme — the aim here being, not the development of an optimum compensation scheme per se, but rather to show that the image models developed earlier *may* be used to develop a scheme which is *capable* of reducing tilt related errors.

However, before the above are addressed the main tool required for these investigations will first be introduced; that is the classifier itself.

6.1. Supervised statistical classification

Classification is the task of assigning objects to groups, or classes, given sets of object measurements. If the classes are known beforehand then the process is termed *supervised* classification. In the context of texture classification the process becomes one of assigning pixels, or groups of pixels, to texture classes, where the sets of "object measurements" are feature vectors comprising features such as Laws' energy masks.

Previous chapters have reviewed and selected three sets of feature measures for use here. These features are however of little use without a method of developing a set of discrimination rules which may be used to assign pixels to texture classes. Hence a simple statistical classifier has been selected. Such classifiers are relatively straightforward to understand and implement [James85] [Tou74], offer reasonable performance [Linnett91a] [Clarke92], and had the advantage of being available to the author. The next section introduces the theory behind these classifiers.

6.1.1. Discriminant theory

Bayes' rule provides the basis for probabilistic classifiers that seek to minimise the "total error of classification" or TEC [James85] [Tou74]. It may be expressed as follows:

> *Assign the pixel with feature vector* **f** *to group* G_i *for which*

$$P(G_i|\mathbf{f}) > P(G_j|\mathbf{f}) \quad \forall j \neq i \tag{6.1}$$

where

> $P(G_i|\mathbf{f})$ is the conditional probability that the pixel with feature vector **f** belongs to group G_i.

Unfortunately these conditional probabilities are difficult to obtain. Bayes' theorem however, expresses them in terms of more easily obtained data :

$$P(G_i|\mathbf{f}) = \frac{P(\mathbf{f}|G_i)P(G_i)}{\sum_{\text{all } i} P(\mathbf{f}|G_i)P(G_i)} \tag{6.2}$$

Thus a maximum likelihood classification rule may be expressed in terms of conditional probabilities, where $P(\mathbf{f}|G_i)$ is the probability of a pixel from group G_i having a feature vector of **f**, and $P(G_i)$ is the *a priori* probability of a pixel belonging to group G_i. To further simplify the classification rule, the associated probability distribution functions are often assumed to be multivariate normal, that is :

$$P(\mathbf{f}|G_i) = \frac{1}{(2\pi)^{n/2}|\mathbf{C}_i|^{1/2}} \exp\left[-\frac{1}{2}(\mathbf{f} - {}_i)' \mathbf{C}_i^{-1}(\mathbf{f} - {}_i)\right] \tag{6.3}$$

where :

> n is the number of feature measures contained within the column feature vector **f**,
>
> $\mathbf{C}_i$ is the n by n variance/covariance matrix of group i,
>
> μ_i is the n element column vector of feature measure means for group i.
>
> $(\mathbf{f} - \mu_i)'$ is the transpose of $(\mathbf{f} - \mu_i)$

Substituting (6.2) and (6.3) into (6.1), taking natural logs (ln), and reversing the inequality [James85, p20] gives the following rule :

assign the pixel with feature vector **f** *to group* G_i *if*

$$\begin{aligned}&\ln\left|\mathbf{C}_i\right| + \left(\mathbf{f} - \;_i\right)' \mathbf{C}_i^{-1}\left(\mathbf{f} - \;_i\right) - 2\ln(P(G_i)) < \\ &\ln\left|\mathbf{C}_j\right| + \left(\mathbf{f} - \;_j\right)' \mathbf{C}_j^{-1}\left(\mathbf{f} - \;_j\right) - 2\ln(P(G_j)) \;\; \forall j \neq i\end{aligned} \tag{6.4}$$

For convenience the terms in the LHS of (6.4), with the exception of the *a priori* probability, are often collected together in one function $d_i(\mathbf{f})$ referred to as the *discriminant function.* Where

$$d_i(\mathbf{f}) = \ln\left|\mathbf{C}_i\right| + \left(\mathbf{f} - \;_i\right)' \mathbf{C}_i^{-1}\left(\mathbf{f} - \;_i\right) \tag{6.5}$$

expanding (6.5) gives

$$d_i(\mathbf{f}) = \ln\left|\mathbf{C}_i\right| + \mathbf{f}'\mathbf{C}_i^{-1}\mathbf{f} - 2\mathbf{f}'\mathbf{C}_i^{-1}\mu_i + \mu_i{}' \mathbf{C}_i^{-1}\mu_i \tag{6.6}$$

This form is known as a *quadratic discriminant* (due to the $\mathbf{f}'\mathbf{C}_i^{-1}\mathbf{f}$ term). If the variance/covariance matrices of all classes are identical then the quadratic and natural log. terms may be eliminated to give a *linear discriminant*

$$d_i(\mathbf{f}) = \mu_i{}' \mathbf{C}_i^{-1}\mu_i - 2\mathbf{f}'\mathbf{C}_i^{-1}\mu_i \tag{6.7}$$

Assuming equal *a priori* probabilities the classification rule now becomes :

assign the pixel with feature vector **f** *to the group* G_i *with the lowest discriminant score* $d_i(\mathbf{f})$

The simpler linear discriminant is used here as it is straightforward to implement and because of its reported robustness and performance [James85]. It assumes a multivariate normal distribution and identical variance/covariance matrices $\mathbf{C}_i$. As these matrices are not normally *identical* they are often replaced by the pooled variance/covariance matrix $\mathbf{C}_p$, in which each element is the average of the corresponding elements of the individual group variance/covariance matrices $\mathbf{C}_i$ [James85].

6.1.2. Supervised classification of test textures

Having decided upon (i) the form of the discriminant function and (ii) the feature set to be used, implementation of a classifier is straightforward. First the training set must be selected, comprising *representative* samples of each texture class. Second, feature images of each sample are generated using the chosen feature set. Third, statistics μ_i and $\mathbf{C}_i$ of the feature image set of each group G_i must be calculated and used to implement the discriminant functions $d_i(\mathbf{f})$. Finally the discriminant functions are built into the classifier as shown in figure 6.1.

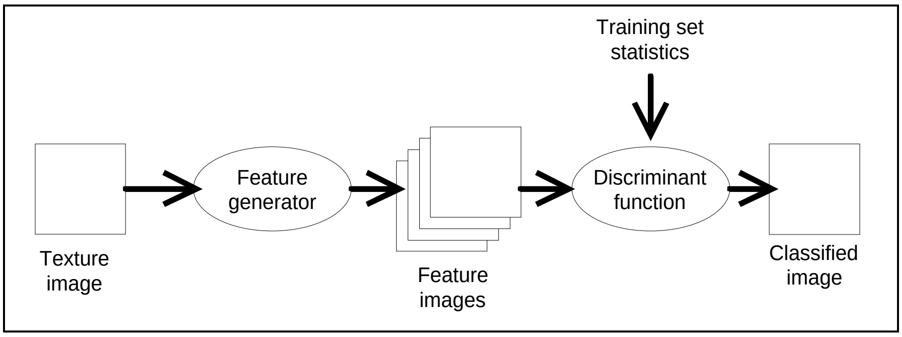

Figure 6.1 - Supervised statistical classification of image texture

To perform a classification of a multi-texture image, feature images are first generated using a feature set such as Laws' energy masks. Secondly, these feature images are used to calculate discriminant scores for each group at each pixel position. The output image resulting from this process is a class map in which the value of each pixel corresponds to the group with the lowest discriminant score at that pixel position. Figure 6.2 illustrates the effect of applying such a classifier to *montage1* — a montage assembled from one directional and three isotropic textures.

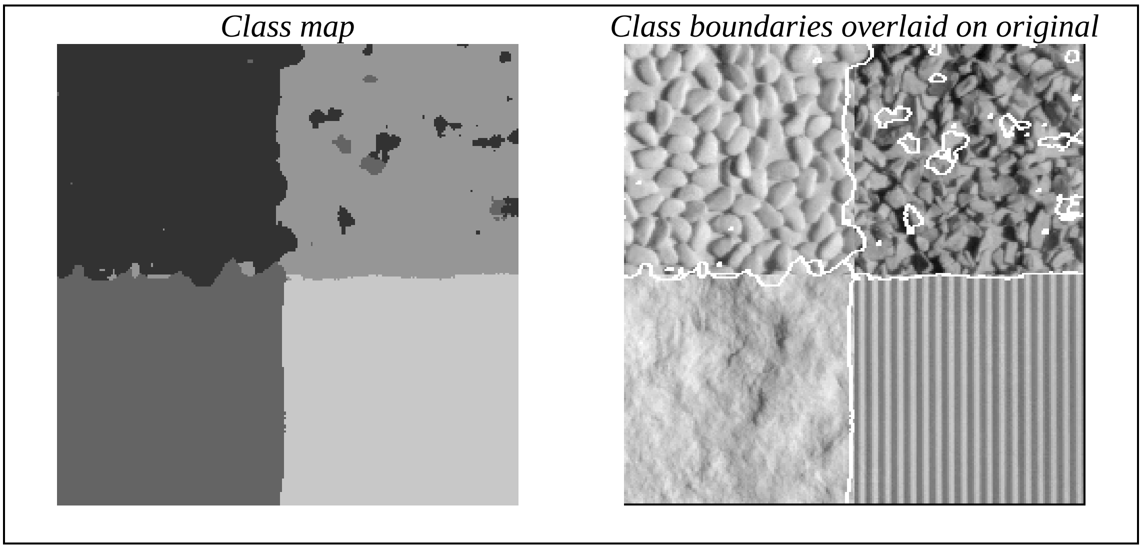

Figure 6.2 - Classification of the four texture image "montage1"

The results were obtained using the Laws' features described in chapter 5. This classifier is referred to here as "*Laws1*" and is defined in table 6.2. It was trained and tested on the image shown in figure 6.2. Note that its three isotropic textures represent a deceptively easy classification task — as these textures have very similar directional characteristics when imaged under the same illumination conditions (see chapter 3). Despite this, the results show that the classifier has been reasonably successful; correctly identifying 96% of the pixels.

6.2. The effect of illuminant variation on classification

If the effect of illuminant variation on a feature set is significant, then it is reasonable to expect that a classifier using such a feature set would be able to discriminate between differing illumination conditions. Hence this section first examines the ability of the *Laws1* classifier to classify images of the same physical texture imaged under *two* lighting conditions, as belonging to *different* classes. The second and third sections directly investigate the effect of illuminant slant and tilt angle variation on classification accuracy.

6.2.1. Discrimination between illumination conditions

In order to test the ability of a statistical classifier to discriminate between differing lighting conditions, a test image "*montage2"* was constructed from four samples of image texture. The four samples consisted of images of *beans1* and *rock1* captured with illuminant tilt angles of 0° and 90°. This test set was used both for training and testing the *Laws1* classifier. Figure 6.3 and table 6.1 contain the results of this classification test. They show that a standard classifier, using Laws' features, is capable of discriminating between different illumination conditions. Thus such features may be of value for the estimation of illuminant tilt. More importantly for this thesis however, is that they show without doubt that the classifier is affected by illuminant variation.

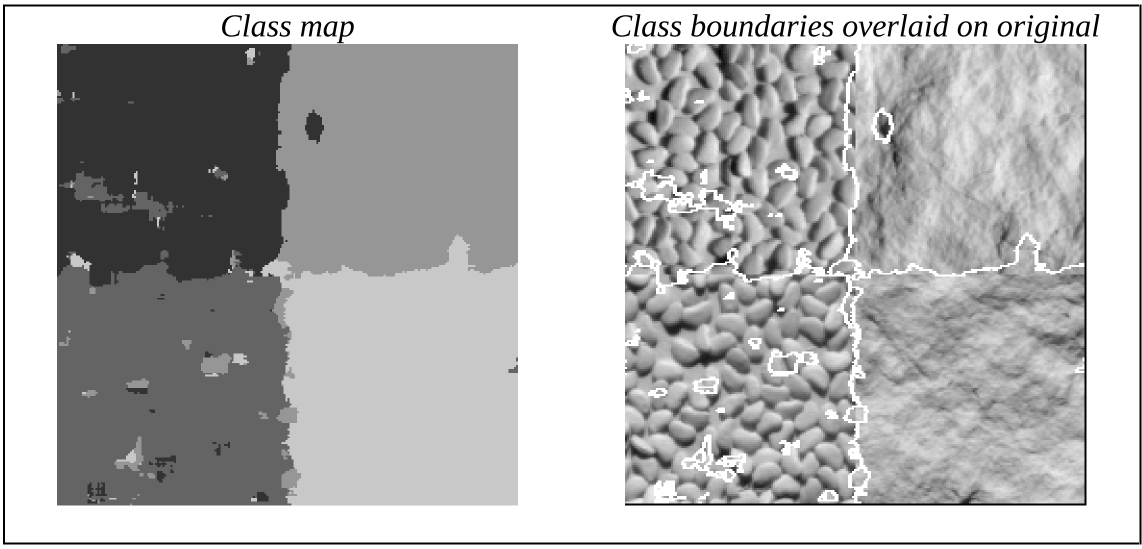

Figure 6.3 - Classification of "montage2" : two physical textures imaged under two illumination conditions

TEC	*beans1*, $\tau = 0°$ (upper left)	*beans1*, $\tau = 90°$ (lower left)	*rock1*, $\tau = 0°$ (upper right)	*rock1*, $\tau = 90°$ (lower right)
5.5%	2.4%	1.8%	1.0%	0.30%

Table 6.1 - Classification errors for figure 6.3

6.2.2. Slant response

The significance of the effects of illuminant variation can only be judged with respect to a *particular* classification task. Classification of textures that differ greatly from one another may not be affected at all, on the other hand textures which are "close" to one another in the feature space may be particularly sensitive to illuminant variation. Here therefore the effect of slant variation on the classification of the test set *montage1* is examined in detail. This test set contains the isotropic textures *beans1, chips1, rock1,* and the directional texture *card1.* It was used to investigate the behaviour of three classifiers the feature sets of which are defined in table 6.2. Each of the feature sets has been defined such that they use the same local window size (e.g. Laws' 5x5 masks together with a 29x29 ABSAVE operator uses a local window or context of 33x33). Each of the classifiers was trained on *montage1* with illumination parameters $\tau = 0°$, $\sigma = 50°$. After training, the classifiers were tested with montages constructed from the same physical textures imaged under a range of illuminant slant angles ($\sigma = 10°, 20°, ...80°$). The results for the *Laws1* classifier are shown in figures 6.4 and 6.5.

Classifier	Feature set
Laws1	Laws' 5x5 masks L5E5, E5L5, E5S5, S5E5, L5S5, S5L5, and R5R5 together with a 29x29 ABSAVE (average of absolutes) macro-statistic.
cooc1	Co-occurrence features CON, COR, ENT and ASM using a 33x33 local window, with displacement vectors **d** = (1,0) and (0,1). The number of grey-levels used *Ng* = 16
frac1	One iteration of Linnett's 3x3 operator with λ = 1 for all seven directional masks, followed by a 31x31 ABSAVE macro-statistic.

Table 6.2 - Definition of feature sets

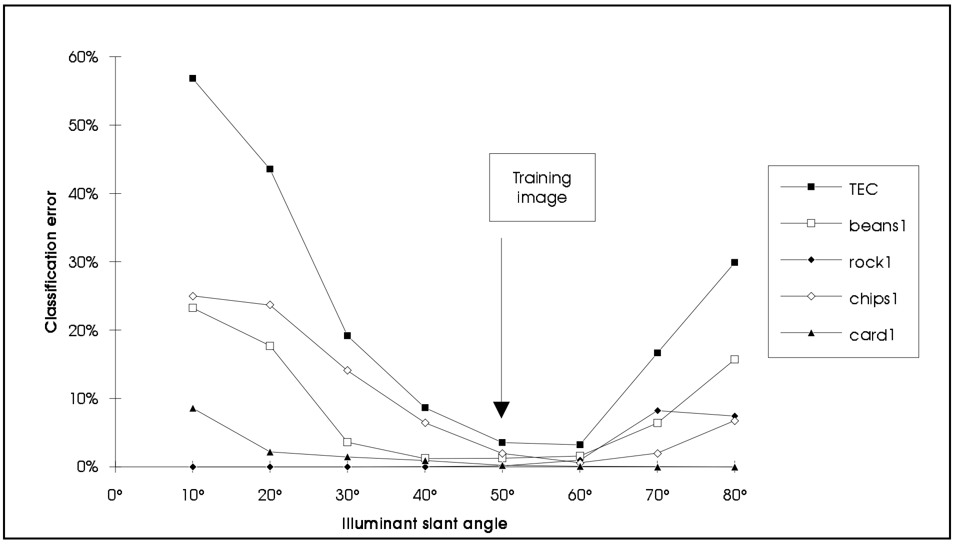

Figure 6.4 - The effect of illuminant slant variation on classifier Laws1

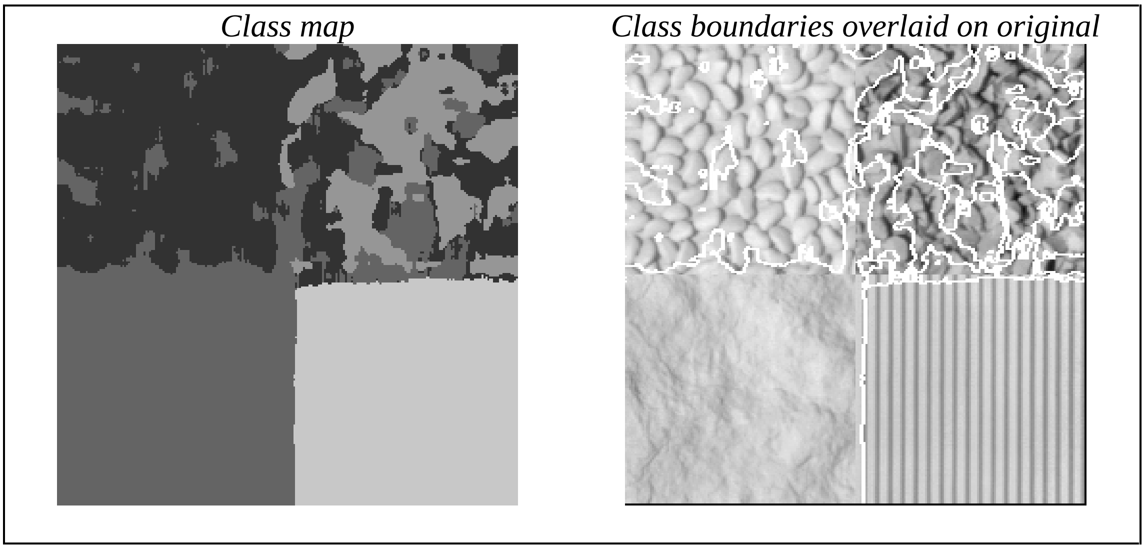

Figure 6.5 - An example of increased failure rate due to variation in illuminant slant (training σ = 50°, test case σ = 30°)

Figures 6.4 and 6.5 show that variation of illuminant slant between training and classification can have a dramatic effect on error rates. Classification using the two other feature sets, *frac1* and *cooc1*, produced similarly catastrophic failures to those shown above. Clearly these classifiers are *not* invariant to changes in the illumination's slant angle.

Normalisation

The image model of topological texture developed in chapters 2 and 3 predicts that normalisation will compensate for slant angle variation. Indeed, chapter 5 showed that normalisation does reduce the variation of the features due to changes in illuminant slant. It is to be expected therefore, that normalisation of images will reduce the error rate of a classifier that has to cope with variation in slant.

Figure 6.6 shows the classification error that results from using *normalised* images (i.e. all images were adjusted to a mean of 127 and a variance of 100 before construction of the montages).

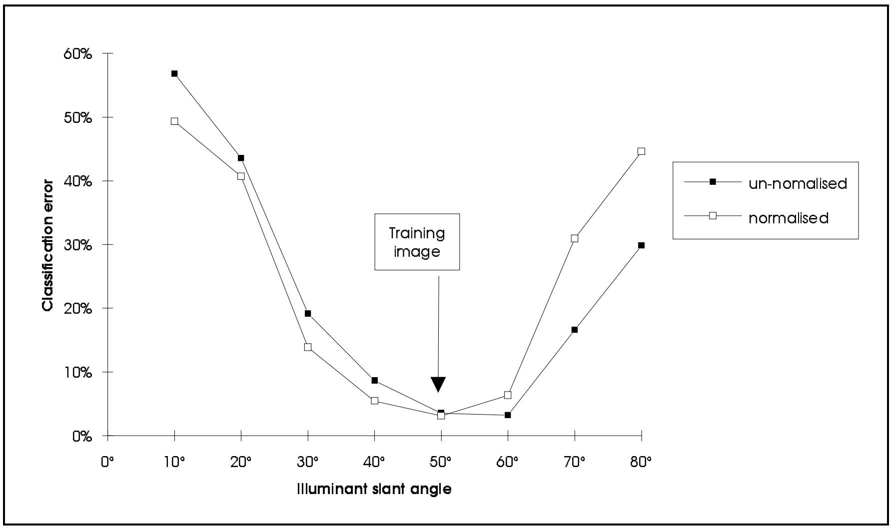

Figure 6.6 - The effect of normalisation on the slant response of the Laws1 classifier (data set : "normalised" montage1)

It can be seen that the use of normalised images produces disappointing results. Although the classification error has been reduced for angles of slant less than 50° it has increased for larger angles. An examination of the slant responses of the features of chapter 5 (figures 5.13, 5.19 and 5.28) shows that in most cases normalisation reduces variation due to changes in slant — especially for angles of 50° or less. Unfortunately it also reduces the separation between class means (see graphs of *beans1*, *chips1,* and *rock1*). This reduction in separation means that the classifiers using normalised images are more sensitive to any changes due to illuminant variation (such as "over-compensation" effects). Hence normalisation of texture images may not necessarily improve a classifier's invariance to illuminant slant.

Tests with *cooc1* and *frac1* classifiers produced similar error rates to those shown above, reinforcing the proposition that normalisation does not necessarily improve a classifier's ability to cope with variation in illuminant slant. Hence the conclusion of this section is that classifiers using feature sets similar to those tested are *not* invariant to changes in illuminant slant whether or not image normalisation is employed.

6.2.3. Tilt response

Chapter 5 showed that Laws', co-occurrence, and Linnett's features, are affected by variation in illuminant tilt. In addition a previous section of this chapter has shown that the *Laws1* classifier is capable of distinguishing between images of the same physical texture captured under differing values of illuminant tilt. Hence it is to be expected that illuminant tilt variation may cause significant problems for supervised texture classification. As with the previous investigation into the effects of illuminant slant, sensitivity to illuminant tilt may only be assessed with respect to a particular classification task. Here therefore, the same test set *montage1* is used as in the previous section. The three classifiers were again trained on textures captured with $\sigma = 50°$ and $\tau = 0°$, but for this experiment the tilt angle (τ) of the test sets was varied in 10° steps from 0° to 180°, while the slant was kept constant at $\sigma = 50°$. The resulting classification error rates are shown below for the *Laws1* classifier (figure 6.7) together with images of one of the worst classifications (figure 6.8).

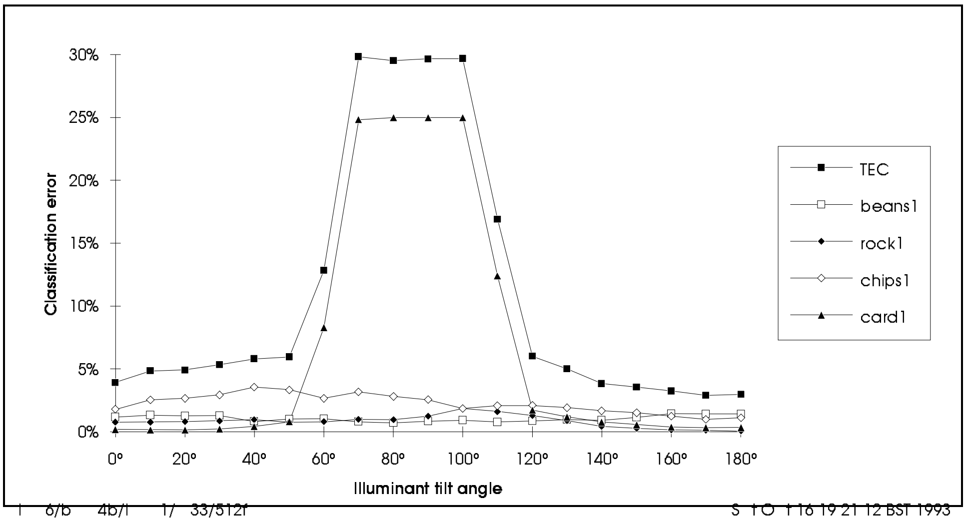

Figure 6.7 - The effect of tilt variation on the classifier Laws1 (data set : montage1)

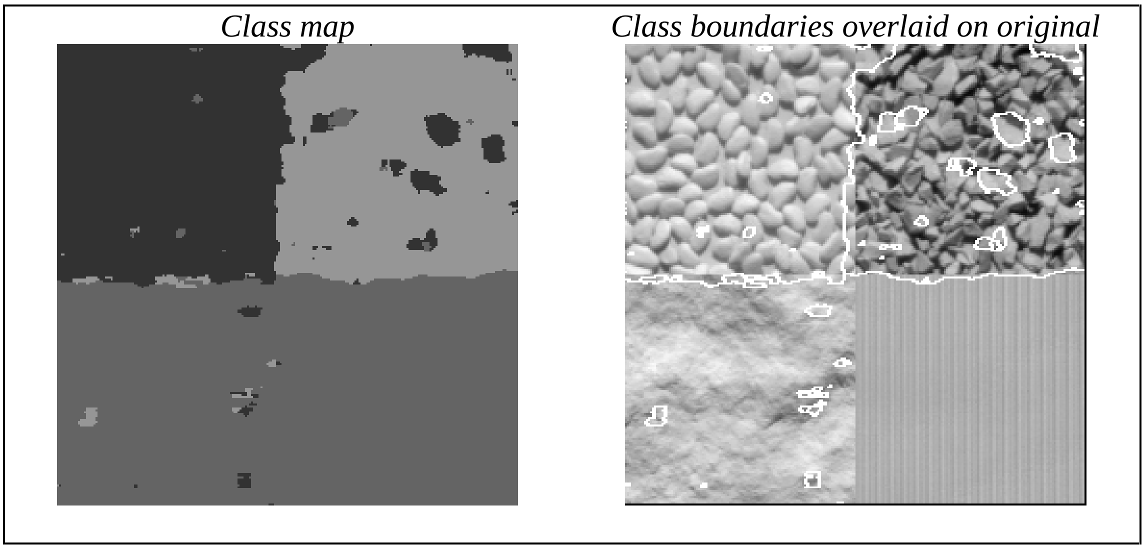

Figure 6.8 - Classification failure at τ = 90° for the Laws1 classifier (data set : montage1)

Figures 6.7 and 6.8 show that, for this data set, the *Laws1* classifier is (i) significantly affected by variation of illuminant tilt, and (ii) that the TEC (total error of classification) is dominated by the failure to correctly classify the majority of class *card1* between tilts of 50° and 120°. Experiments on the *cooc1* and *frac1* classifiers using the same data set gave similar results (see figure 6.9).

These results clearly demonstrate that variation of illuminant tilt between training and classification sessions can have a dramatic effect on the accuracy of a statistical classifier.

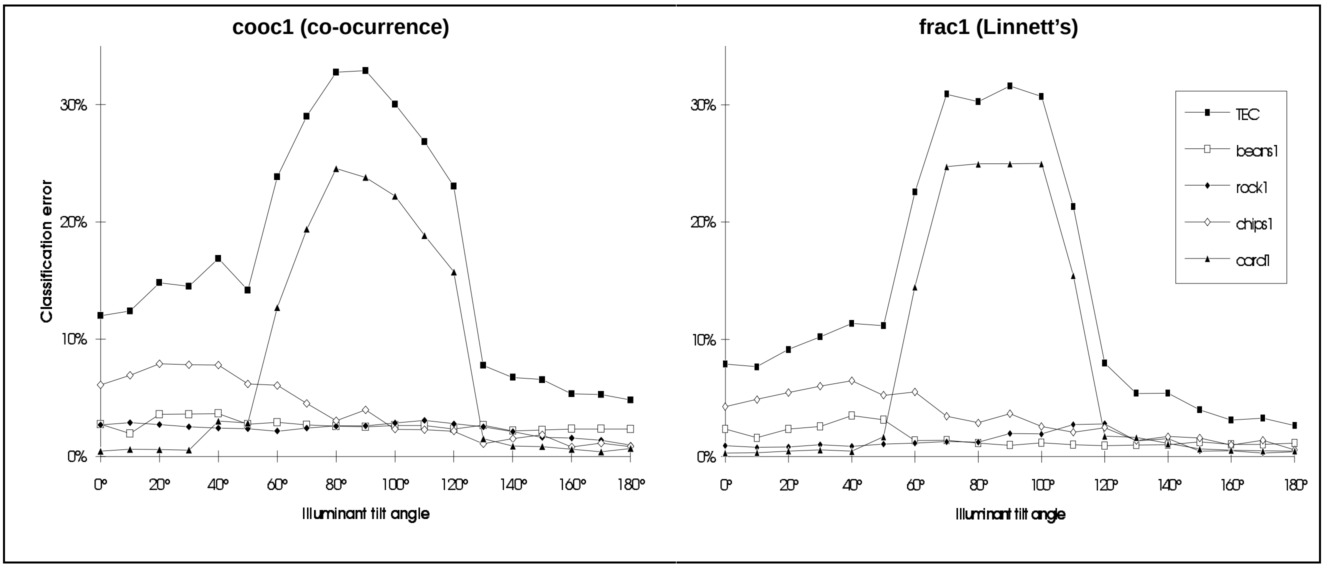

Figure 6.9 - The effect of illuminant tilt variation on "cooc1" and "frac1" classifiers

6.2.3.1 Normalisation

In chapter 5 it was shown that normalisation does not have a significant effect on images of *isotropic* texture taken under varying values of illuminant tilt — as although a change in tilt does alter the balance between the texture energy in differing directions, it does not alter the *overall* energy of the texture image. Thus normalisation has the same effect on each image of an isotropic texture regardless of illuminant tilt. However, the same was shown not to be the case for *unidirectional* textures. The variance of an image of a unidirectional texture *does* vary with illuminant tilt. That is as τ approaches 90° to the texture direction, image variance is reduced. Normalisation however, makes the variance of each image identical regardless of tilt. Thus, in theory, normalisation should reduce the effect of tilt variation on images of unidirectional texture. Figure 6.10 and 6.11 illustrate the effect of applying image normalisation. Note that each texture was normalised *before* being added to the test montage — simulating an ideal local normalisation process.

When figure 6.10 is compared with the un-normalised error rates (figure 6.7) it is clear that normalisation has reduced the mis-classification of the directional texture *card1*, and hence it has also reduced the TEC (total error of classification).

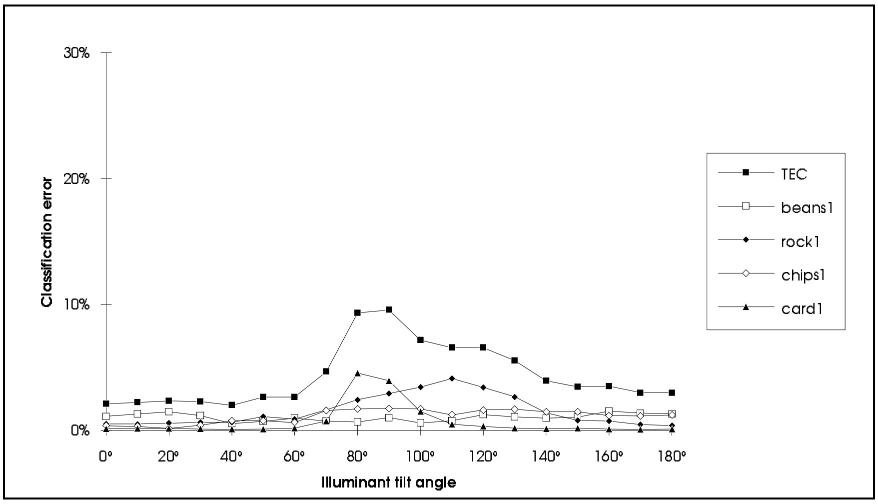

*Figure 6.10 - The effect of normalisation on the previous classification problem (Laws1 classifier; data set : **normalised** montage1)*

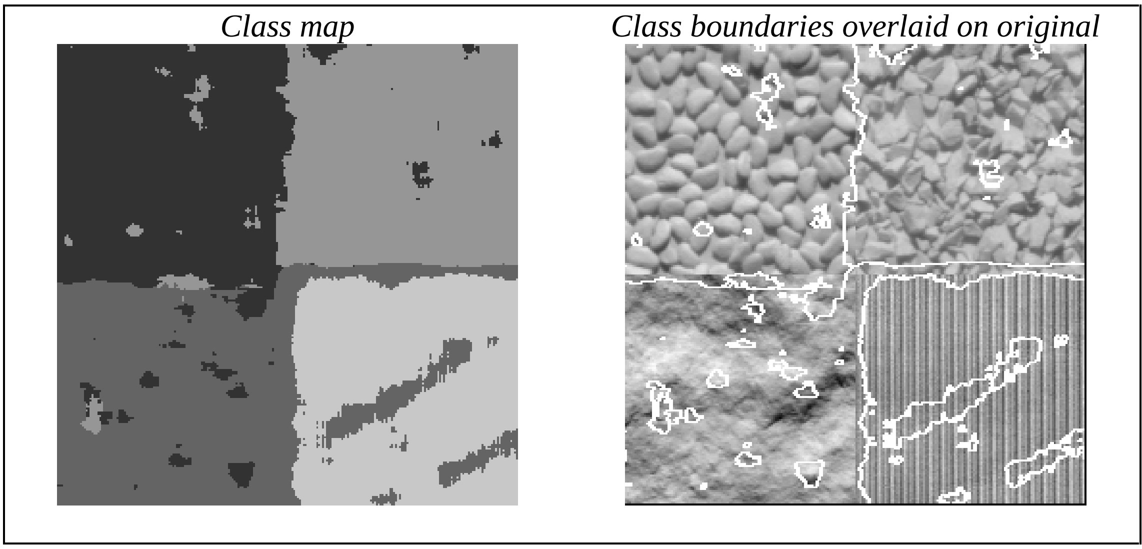

*Figure 6.11 - Classification at $\tau = 90°$ (Laws1 classifier, **normalised** montage1)*

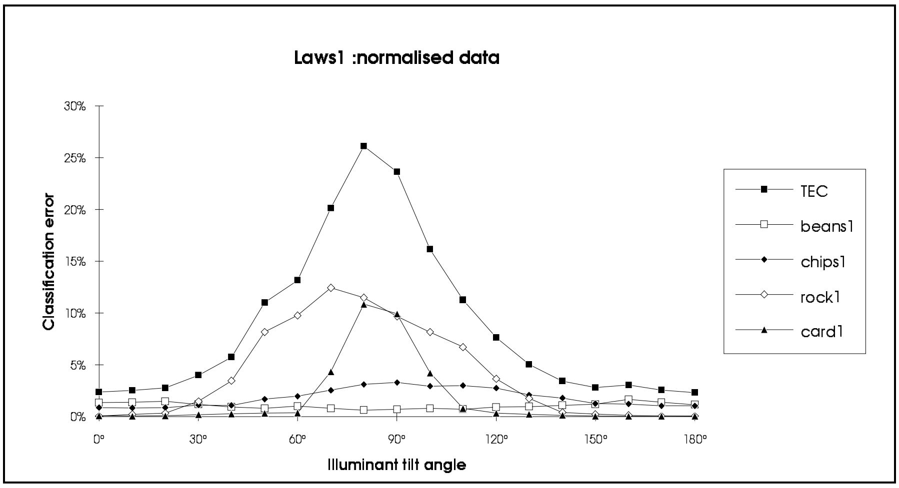

Figure 6.12 - The effect of tilt variation on the Laws1 classifier using normalised images (data set : montage3)

However, because a texture's variance is one of its distinguishing characteristics, normalisation might also be expected to *reduce* the classification accuracy in some cases. Hence another test set, *montage3,* was constructed using different samples of *rock1* and *chips1*. It was presented to the *Laws1* classifier as before. The resulting error rates are displayed in figure 6.12. It shows that while normalisation has reduced the classification error of the directional texture *card1* in *montage3*, it has also unfortunately *increased the error* associated with the isotropic texture *rock1*. Thus normalisation may actually increase error rates, as well as decrease them.

Figure 6.13 shows the results of repeat experiments for the *cooc1* and *frac1* classifiers (co-occurrence and Linnett's features respectively). Again the graphs show that the errors associated with the directional texture *card1* have been significantly reduced, and that those associated with the isotropic textures have increased particularly those of *beans1*. Re-examination of table 5.1 reveals further supporting evidence that normalisation can reduce classification accuracy — the class separation figures of all of Linnett's features are significantly lower in their normalised form. The same holds for co-occurrence features with the exception of the *COR* measure.

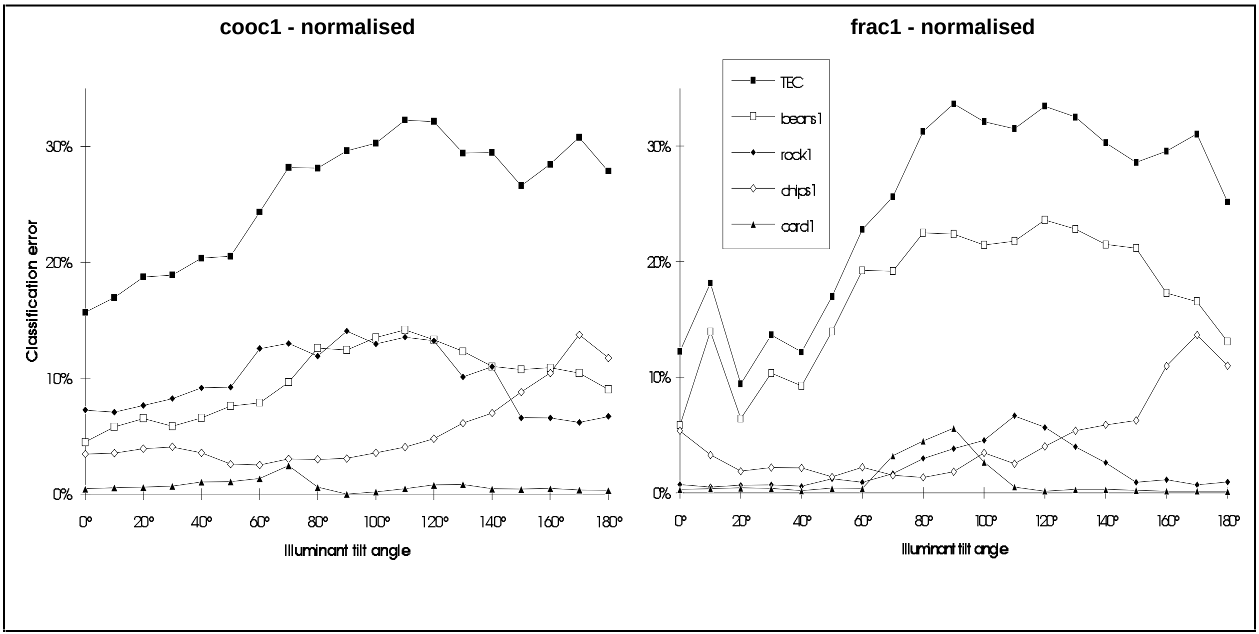

Figure 6.13 - The effects of image normalisation on cooc1 and frac1 classifiers (data set : montage1)

Thus although normalisation does reduce the classification error rates of the directional class *card1*, it has also been shown to *increase* the error rates of some of the isotropic textures.

6.2.4. Summary of illuminant variation investigation

This section has described the effects of variations in the illuminant's slant and tilt angles on the classification of directional and isotropic textures. The illuminant tilt and slant responses of three classifiers have been presented, and the effects of normalisation have also been investigated. To summarise :

- A classifier using Laws' features has been shown to be capable of discriminating between two sets of illumination conditions.
- Variation in illuminant tilt and slant have both been shown to significantly reduce the classification accuracy of three classifiers when applied to a test montage of isotropic and directional textures.
- Image normalisation was shown to have little effect on these slant induced errors.
- Image normalisation was shown to *reduce* the tilt related classification errors of the directional texture *card1*.
- Image normalisation was shown to *increase* the tilt related classification errors of some of the isotropic test textures.

6.3. Compensation for illuminant tilt variation

The previous section has shown (i) that variation in illuminant tilt can significantly affect supervised classification of three-dimensional texture, (ii) that normalisation can help compensate for such variations where directional textures are concerned, and (iii) that normalisation may actually degrade a classifier's ability to classify isotropic texture. Normalisation is however only one of a number of possible compensation schemes. Some alternatives are now proposed.

Proposal 1

The simplest solution is to train the classifier over the range of illuminant conditions that are likely to occur during classification sessions. However, such an approach may significantly reduce the accuracy of the system from its potential optimum, as the variance of each class is likely to be higher than would be the case for fixed illumination conditions.

Proposal 2

If the tilt angle of the illuminant is varied during training, then a family of discriminant functions may be developed. This would provide what is essentially a lookup table of discriminants indexed by tilt angle. Alternatively the tilt angle may be used as a feature itself. Both approaches are only viable if the appropriate training sets are available and the tilt angle of the test data is known. They would also require additional resources for gathering and handling the data and performing the training. It must be said however, that such methods are likely to be simple and may well produce good results.

Proposal 3

Use unsupervised classification techniques; e.g. *k-means* clustering [Tou74] for training set identification followed by a statistical classifier [Linnett91a]. If illuminant variation affects each texture in a similar manner then a change in lighting will impart approximately the same displacement in feature space to each texture class. Thus it is to be expected that unsupervised techniques will not be as severely affected as supervised ones — as the clustering will *track* the changes in class centres and the discriminants will be adjusted accordingly. However, the segmentation of an image into homogeneous regions would not at first sight be of great benefit if the job is to classify textures into previously defined groups. Nevertheless this segmentation process is of value — as larger texture regions maybe used for more sophisticated feature generation processes. Such an approach may for instance enable FFTs to be used to provide information on the radial shape of a texture's PSD — as in chapter 2 it was suggested that such characteristics are *intrinsic* to the texture. Thus segmentation using unsupervised techniques followed by extraction of features based on PSD radial shape may provide a lighting invariant classification scheme.

Proposal 4

Reverse the *directional filtering* effect of single point illumination by using a family of compensating filters each one constructed for a particular value of tilt. Thus each image, be it a training or test image, would be passed through a filter corresponding to the illuminant tilt angle under which the image was captured. Hence in theory tilt related characteristics would be removed — allowing classifiers to be trained under one set of illumination conditions but to be used with arbitrary tilt angles. Note that this method requires the illuminant's tilt angle either to be known or obtainable from a reliable estimator.

Choice

It is not practical within this thesis to address all of the above avenues. It was therefore decided to choose just one for investigation here. The first two proposals are straightforward. However, the first is unlikely to provide good performances for difficult texture classification tasks and the second requires extensive training. The third proposal is interesting in that it does not need illuminant tilt as input, but it is more speculative and would be more complex to implement than the other proposals. The last proposal does require illuminant tilt to be known, but is simple to develop, and offers the potential advantage that training requirements would be significantly reduced compared with proposal 2. For these reasons the fourth proposal based on the development of a compensating filter will be investigated.

6.4. Frequency domain tilt-compensation

This section proposes a tilt-compensation method which is based upon the frequency domain model of image texture developed earlier. Its purpose is not to develop an optimum compensation scheme and extensively test it; rather it is to show that the model of image texture *may* be used to develop a scheme which is *capable* of reducing tilt related errors.

Chapters 2 and 3 developed an image model of topological texture. If slant (σ) is constant then, as in chapter 5, the model reduces to equation (5.13) :

$$F_I(\omega,\theta) = F_S(\omega,\theta).F_\tau(\omega,\theta).k_\sigma$$

and substituting (3.8) gives

$$F_I(\omega,\theta) = F_S(\omega,\theta).(m_\tau \cos(\theta-\tau)+b_\tau).k_\sigma \qquad (6.8)$$

Thus if the illuminant tilt is known, a *tilt-compensation filter* of the form

$$H_{\tau c}(\omega,\theta) = \frac{1}{m_\tau \cos(\theta-\tau)+b_\tau} \qquad (6.9)$$

may, in theory, be applied to remove variations due to changes in tilt angle. This filter must of course be applied to all test images *and* to all training images. It should be applied before feature generation; as shown in figure 6.14.

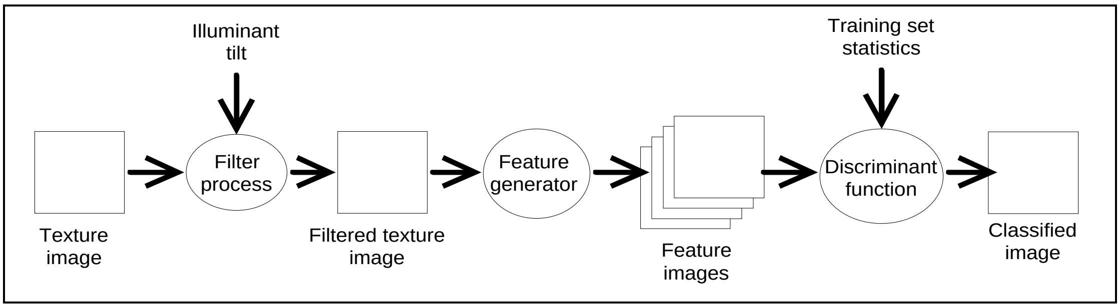

Figure 6.14 - The use of a tilt-compensating filter in the texture classification process

Hence the main advantage of this scheme is that training images only need to be obtained under a single set of illuminant tilt conditions — as tilt-compensation filters will in theory compensate for any variations due to changes in τ.

The coefficients m_τ and b_τ in (6.9) were obtained in the first instance by taking an average of estimates derived from four isotropic textures. The estimates were calculated by using a least squares fit of the tilt response model (i.e. the inverse of equation 6.9) to a set of polar plots of the two-dimensional magnitude spectra. These plots which were normalised to have a mean = 1.0, were of the textures *rock1*, *beans1*, *chips1*, and *stones1*, imaged with $\tau = 0°$. (Figure 3.26 shows un-normalised polar plots of the four textures.) Thus a value of 0.6 was used for both m_τ and b_τ. The resulting family of tilt-compensation filters, referred to as "*F1*" in the following text, is defined below.

$$H_{F1}(\omega, \theta) = \frac{1}{0.6\cos(\theta - \tau) + 0.6} \tag{6.10}$$

Figure 6.15 shows the magnitude frequency response of filter *F1*(τ = 0°).

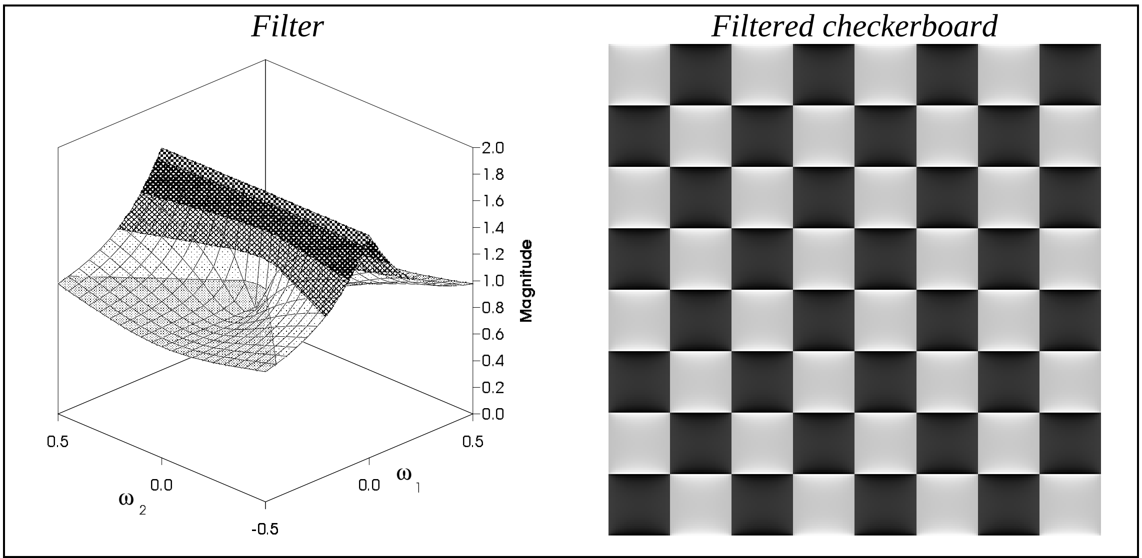

Figure 6.15 - Magnitude frequency response of F1(τ = 0°), and its effect on a checkerboard image.

The τ = 0° filter amplifies components with an angle θ = 90° and attenuates those with θ =0°. This effect is readily apparent in the image, shown in figure 6.15, which results from the application of *F1*(τ = 0°) to a checkerboard image.

Unfortunately application of this filter family to images of a test set, comprising isotropic textures (*set a*), actually *increases* the average tilt sensitivity of the Laws' features (see table 6.3). A closer examination reveals that only the higher frequency masks R5R5, E5L5, and E5S5 were adversely affected, which suggests that the tilt response model above is inadequate at higher frequencies.

6.4.1. An improved frequency domain model

The previous section found that application of the tilt-compensation filter family *F1* can actually *increase* the tilt sensitivity of texture features, rather than decrease them as intended. Hence this section investigates the magnitude spectra of the four isotropic textures in more detail — the aim being to provide a frequency domain model which will facilitate the development of a tilt-compensation filter that does reduce tilt sensitivity.

It was noted in the previous section that the higher frequency feature measures were adversely affected compared with their lower frequency counterparts. Here therefore, the polar characteristics of texture magnitude spectra are examined over a number of frequency bands. This contrasts previous polar plots in which the magnitude response was averaged over the whole radial frequency range for each value of θ. Thus each of the plots on the graph below shows the polar characteristics of one of a series of concentric rings taken from the two dimensional magnitude spectrum of *rock1*. Each plot is labelled with the centre frequency of the "ring".

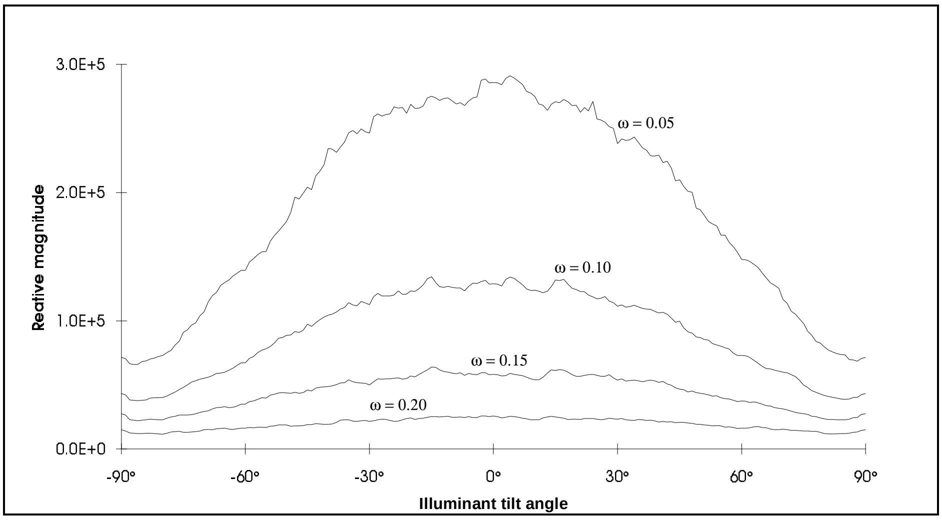

Figure 6.16 - Polar frequency characteristics of rock1 texture (τ = 0°)

From the above it can be seen that energy in the texture *rock1* falls off with frequency. It is assumed that this is a function of the topological texture and will therefore be ignored here. Thus for the purposes of developing a tilt-compensation filter, polar plots are normalised to have a mean = 1.0. Figure 6.17 shows the result of plotting these normalised values against cos(τ - θ) for two values of frequency (ω). From this graph it can be seen that there is an approximate linear relationship with cos(θ - τ) at both frequencies, but that the values of the linear coefficients change with frequency.

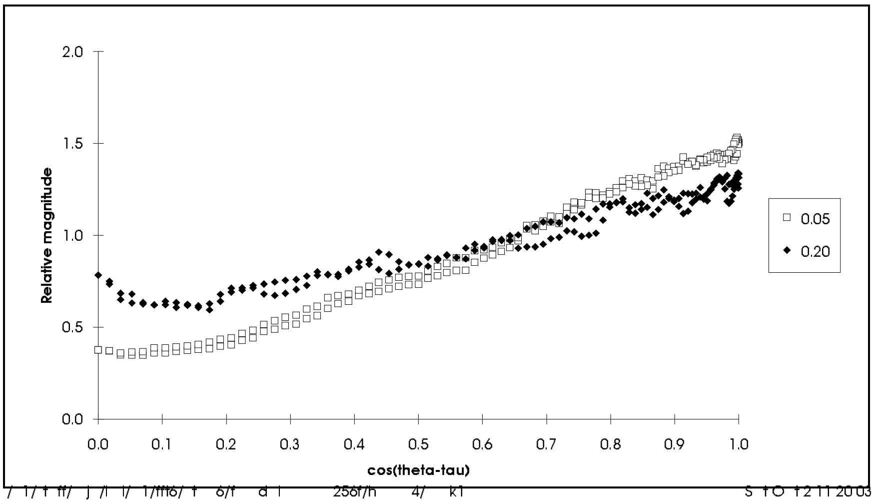

Figure 6.17 - Plot illustrating the $F_I(\omega,\theta) \propto m_\tau cos(\theta - \tau) + b_\tau$ relationship for $\omega = 0.05$ and 0.20 times the sampling frequency

Figures 6.18 and 6.19 provide a more extensive view of the behaviour of these coefficients as a function of frequency.

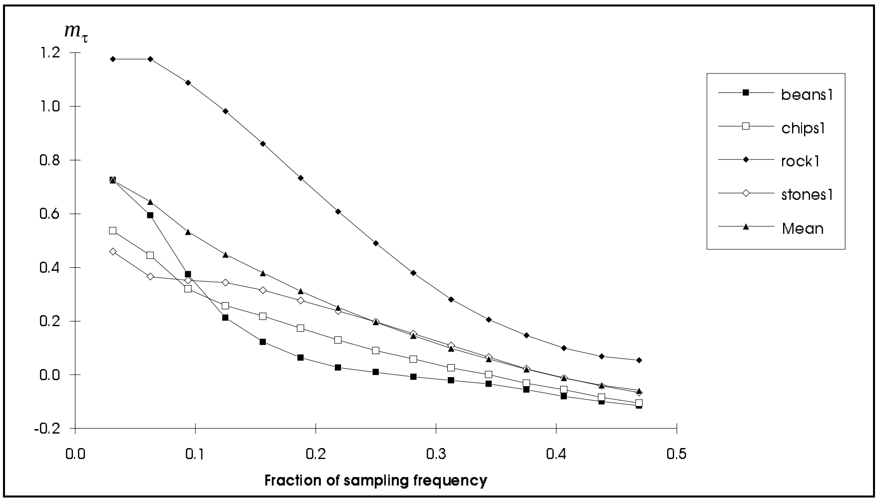

Figure 6.18 - Variation of m_τ with frequency

These estimates of m_τ and b_τ were obtained by averaging least squares estimates at $\tau = 0°$ *and* 90° in order to reduce any directional artefacts that might have been introduced by the data capture or analysis processes.

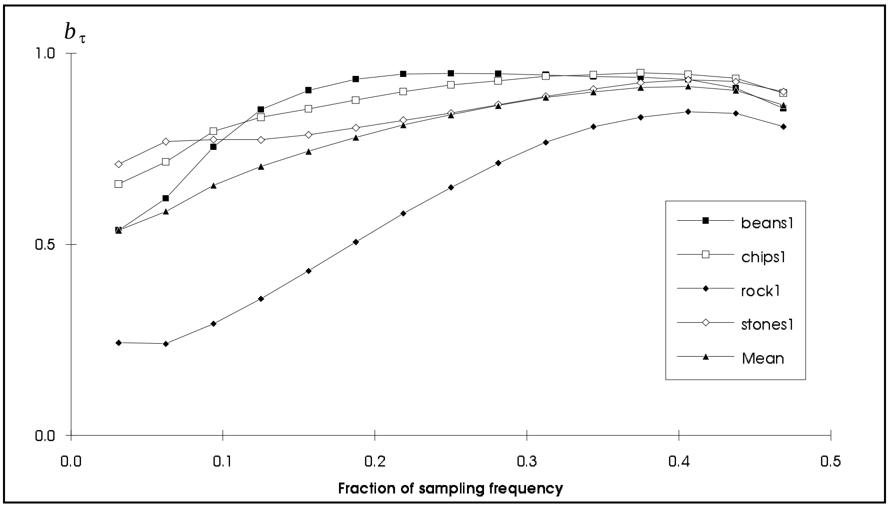

Figure 6.19 - Variation of the parameter b_τ with frequency

What is clear from the above two graphs is that the directional characteristics exhibited are strongest at low frequencies — as the Nyquist frequency is approached the polar plots tend towards a flat, isotropic response (i.e. m_τ = 0, b_τ = 1). One explanation for this behaviour is that, as has been shown in chapter 3, the energy of the textures reduces with increasing frequency. Thus noise will become more significant as frequency increases and hence if the noise is isotropic, it will tend to flatten the polar response.

A simple model was developed in order to account for this behaviour. Least squares estimates of the linear behaviour of m_τ and b_τ as a function of frequency were derived from the mean behaviour of the four textures giving :

$$m_\tau = -1.8\,\omega/\omega_s + 0.7,\;\; b_\tau = 0.8\,\omega/\omega_s + 0.6 \qquad (6.11)$$

where

ω_s is the sampling frequency.

However, (6.11) gives a negative value of m_τ for ω/ω_s > 0.39, that is the directional characteristic of the resulting filter would be the *inverse* of that predicted in chapter 2. Thus the model was modified to a more conservative set of coefficient definitions :

$$\left.\begin{aligned} m_\tau &= -1.4\,\omega/\omega_s + 0.7 \\ b_\tau &= 0.8\,\omega/\omega_s + 0.6 \end{aligned}\right\} 0 < \omega/\omega_s < 0.5$$
$$\left.\begin{aligned} m_\tau &= 0.0 \\ b_\tau &= 1.0 \end{aligned}\right\} \omega/\omega_s \geq 0.5 \qquad (6.12)$$

Note that the filter defined by (6.12) in conjunction with (6.9) has unity gain at frequencies $\omega/\omega_s \geq 0.5$, as opposed to the inverse directional characteristic described above. Hence this modified model (6.12) was used together with equation (6.9) to specify the "*F2*" family of tilt-compensation filters.

6.4.2. Filter implementation

Both *F1* and *F2* filter families were implemented in the frequency domain using forward and inverse FFTs (fast Fourier transforms) as depicted in figure 6.20. First, the two-dimensional magnitude spectrum of the required filter is generated using the illuminant tilt angle as input to the *F1* filter equation (6.10) or the *F2* filter equations (6.9) and (6.12) as required. Second, the texture image is FFTed to provide real and imaginary component images of its complex spectrum. Third, both the real and imaginary images are multiplied, coefficient by coefficient, by the filter image. Finally the filtered real and imaginary images are inverse transformed back into the spatial domain to provide the filtered texture image.

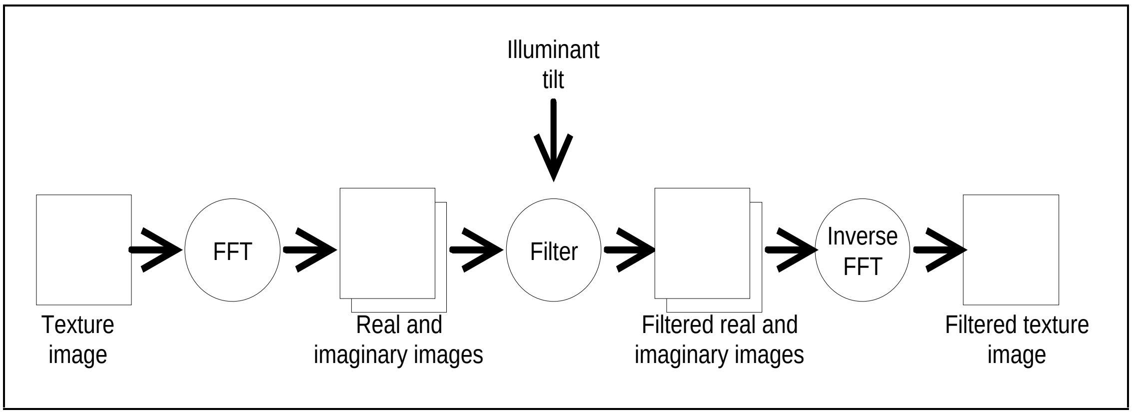

Figure 6.20 - Frequency domain filtering

In comparison with the spectral analysis described in chapter 3, circular Hann windows and the spatial averaging of the Welch periodgram method were not employed. Spatial averaging which was used in chapter 3 purely to aid interpretation is not required here; while artefacts introduced by the forward transform, through the use of non-circular windows, are largely removed by reverse transforms using the same window.

Figure 6.21 shows the result of applying an *F2* filter to an image of the texture *rock1*.

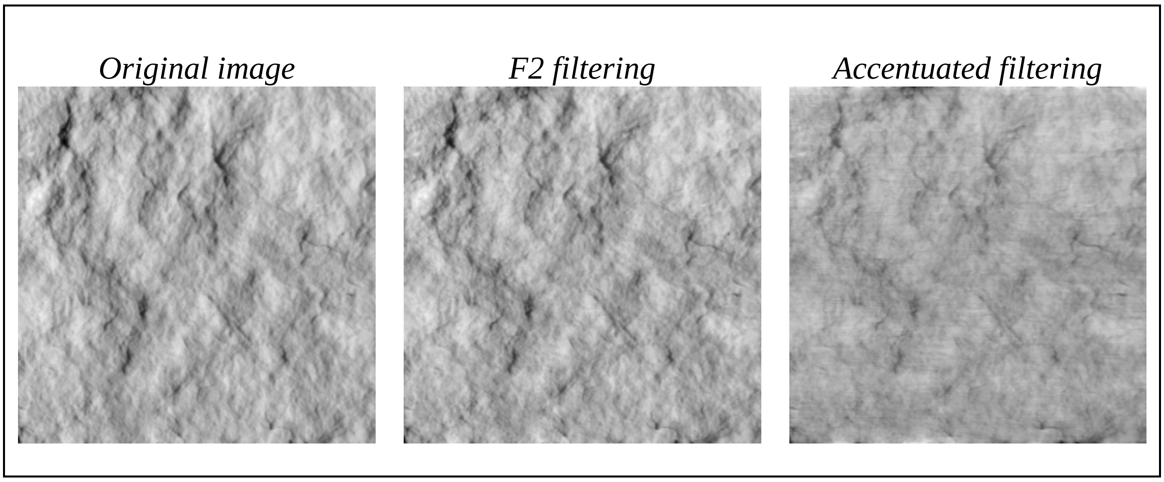

Figure 6.21 - The effect of filter F2(τ = 0°) on the texture "rock1"

As the effect is difficult to discern an *accentuated* version of the filtering is also shown (in which the directional effect has been exaggerated).

6.4.3. Effect of tilt-compensation on features

If the *F1* and *F2* filters are useful for tilt-compensation then their application to texture images will reduce the feature measures' tilt sensitivities. That is the separation between a feature's distributions at differing angles of illuminant tilt should be reduced by the filters.

Figure 6.22 shows the distributions (histograms) of the output of a tilt-compensated Laws' L5E5 texture measure. It has been applied to four image textures : two physical textures each imaged at two values of tilt ($\tau = 0°$ and $90°$).That is each texture image was processed with the appropriate *F2* compensation filter before application of the L5E5 operator.

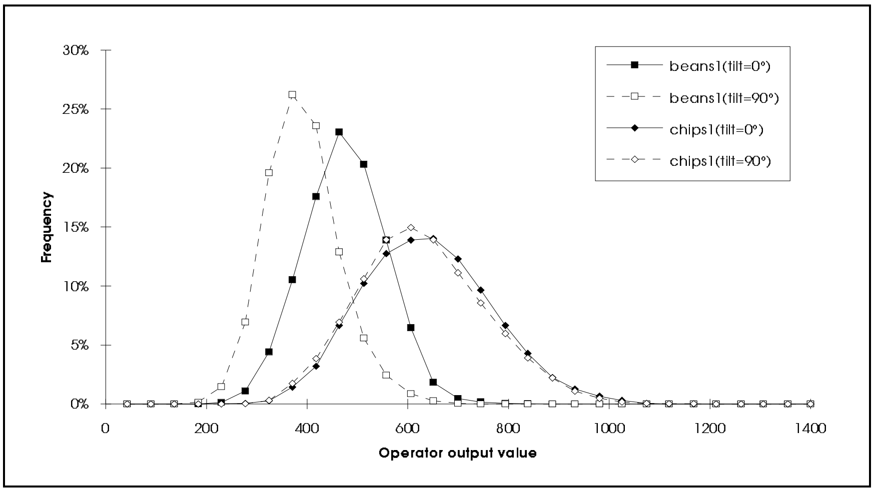

Figure 6.22 - The effect the "F2" filters on tilt behaviour of Laws' L5E5 feature

If these histograms are compared with figure 5.7 (which shows the distributions of the same operator used directly on the original images) it can be seen that the *F2* filters :

(i) have not significantly distorted the shape of the distributions,

(ii) have reduced the displacement of the mean of class *beans1* due to change in τ, and

(iii) have almost eliminated the displacement of the mean of *chips1*.

In addition, if the above graph (figure 6.22) is compared with that showing the result of normalisation (figure 5.11), it can be seen the *F2* filters have not reduced the separation between the class means as has happened for normalisation.

The above is a qualitative, subjective assessment and contrasts with the quantitative objective measure of tilt sensitivity developed in chapter 5. The metric, developed from the Mahalanobis distance, was defined to aid comparison of texture measures. Table 6.3 below contains the results of applying this metric to Laws' texture measures using images pre-processed with the *F1* (512fF1) and *F2* (512fF2) filter sets. Results using the original 512x512 images (512f) are repeated here for convenience. In addition class separation measures are shown to allow the relative effect of the tilt-compensation filters to be assessed. Data sets *set a* and *set b* are as defined in chapter 5. That is *set a* contains only isotropic textures whereas *set b* contains the unidirectional texture *card1* in addition to the textures of *set b.*

			Mean tilt sensitivity		Class separation	
			set a	*set b*	*set a*	*set b*
Mean 512f			1.21	6.31	0.09	0.08
Mean 512N			1.07	12.03	1.92	0.66
Mean 512fF1			1.79	3.13	0.10	0.08
Mean 512fF2			0.78	2.78	0.09	0.07
Laws	*E5E5*	512f	0.16	0.15	0.00	0.00
Laws	*E5S5*		0.41	0.35	0.02	0.02
Laws	*L5E5*		2.55	18.03	0.28	0.21
Laws	*L5S5*		2.90	12.94	0.15	0.14
Laws	*R5R5*		0.05	0.10	0.01	0.00
Laws	*E5E5*	512f F1	1.18	1.13	0.03	0.03
Laws	*E5S5*		1.53	1.64	0.03	0.03
Laws	*L5E5*		1.38	6.83	0.29	0.22
Laws	*L5S5*		1.28	2.59	0.13	0.12
Laws	*R5R5*		3.59	3.45	0.01	0.00
Laws	*E5E5*	512f F2	0.69	0.59	0.01	0.01
Laws	*E5S5*		0.59	0.53	0.02	0.02
Laws	*L5E5*		1.29	8.96	0.29	0.20
Laws	*L5S5*		1.25	3.72	0.14	0.11
Laws	*R5R5*		0.05	0.10	0.01	0.00

Table 6.3 - Tilt sensitivity and class separation of Laws features pre-filtered with F1 and F2 filters. The original floating point figures (512f) are repeated here for convenience.

The tilt sensitivity figures above show that

(i) The filter set *F1* reduces the average tilt sensitivity of Laws' features when used on *set b* (containing a directional texture) but the same filter set *increases* the tilt sensitivity when used with the isotropic data set *set a*,

(ii) The filter set *F2* reduces average tilt sensitivity in *both* cases, and

(iii) neither filter set markedly affects class separation.

Thus these results indicate that pre-processing with the *F2* filter set should reduce tilt related classification errors. It will therefore be used in the next section which describes an investigation into the effect of tilt-compensation on classification error.

6.4.4. Effect of tilt-compensation on classification

This section analyses the effect of the *F2* tilt-compensation filter family on tilt related classification errors. These experiments mirror those described in section 6.2.3, which determined the consequences of varying the illumination's tilt angle on uncompensated[1] images. The same four texture data set (*montage1*) is used here. Training is again performed on images captured at $\tau = 0°$ with classifications being processed at a range of illuminant tilts angles. In this case however, all images are passed through the appropriate *F2* filter (selected by tilt angle τ) before feature processing.

The investigation into this tilt-compensation scheme is reported in three parts : firstly the class error rates are discussed; secondly the total error rates of uncompensated, normalised and tilt-compensated schemes are compared; and thirdly the distribution between isotropic and directional errors is presented.

a) Error rates of individual textures of *montage1*

Figures 6.23 and 6.24 show the results of the first tilt-compensation experiment — performed with the *Laws1* classifier.

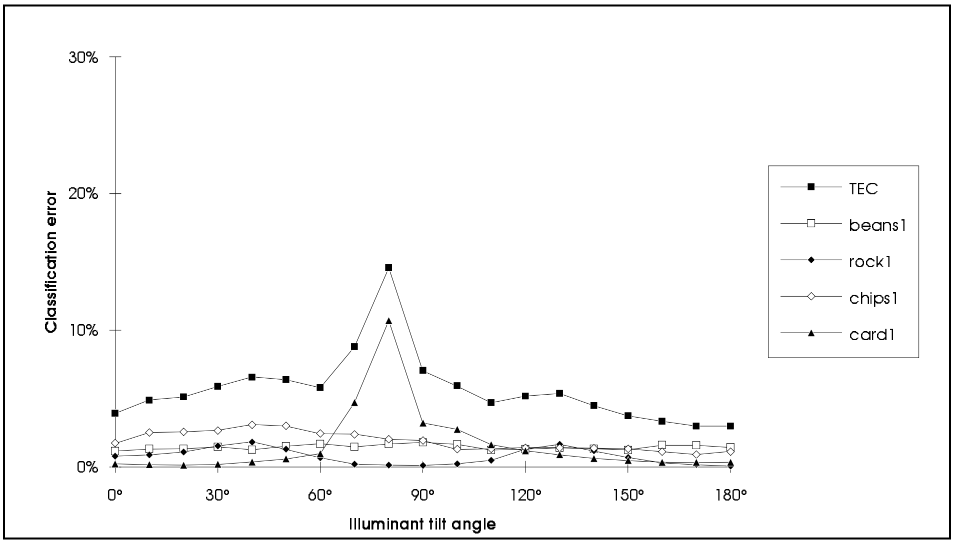

Figure 6.23 - The effect of tilt-compensation on the Laws1 classifier (data set : F2 tilt-compensated montage1)

All images of the data set *montage1* were pre-processed with the appropriate *F2* filter before feature processing. As in previous tilt experiments the illuminant slant angle was maintained at $\sigma = 50°$, the classifier was trained at an illuminant tilt angle $\tau = 0°$, and it

[1]Note that the term "uncompensated" is used in this and subsequent sections to refer to images that have been neither normalised nor pre-processed with a tilt compensation filter.

was tested over the range of tilt angles (τ = 10°, 20° ....180°). Figure 6.24 shows the classification result which occurred at τ = 90°.

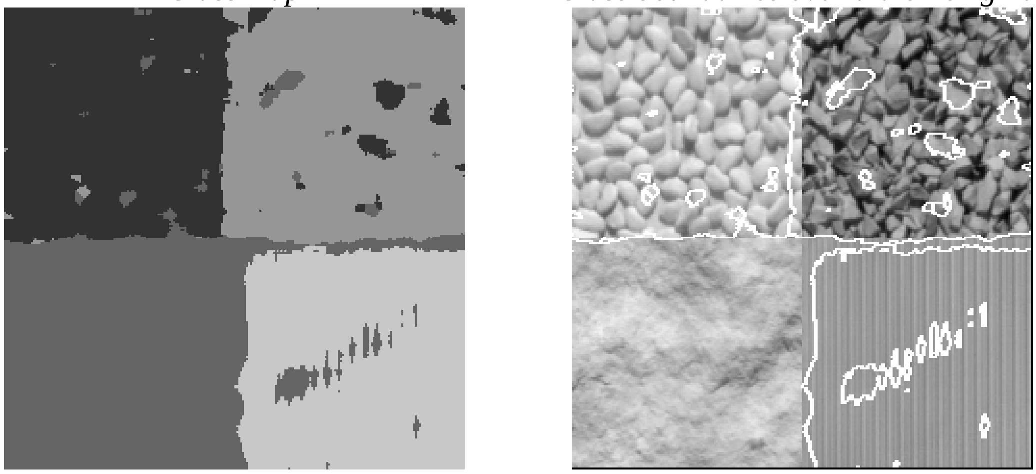

Figure 6.24 - The effect of tilt-compensation on the classification at τ = 90° (Laws1 classifier, data set : F2 tilt-compensated montage1)

Comparison of the above results with the equivalent uncompensated versions (figures 6.7 and 6.8) show that the classification errors associated with the directional texture *card1* have been significantly reduced. This was to be expected, given the reduced tilt sensitivities of the tilt-compensated Laws' features (see table 6.3). The error rates of the isotropic textures however, do not show a similar reduction. The flat graphs of error rates of the uncompensated isotropic textures, shown in figure 6.8, suggest that variation in illuminant tilt does not affect the appearance of these textures enough to cause significant mis-classification. Hence it is not surprising that the tilt-compensation scheme does not reduce isotropic error rates in this instance.

Figure 6.25 below shows the result of using the *F2* filters with co-occurrence and Linnett's features. In the case of the former, the graphs are not as convincing as for the Laws' features — tilt-compensation has reduced the error rate around τ = 90° but has actually increased it at tilt angles of 30° and 40°.

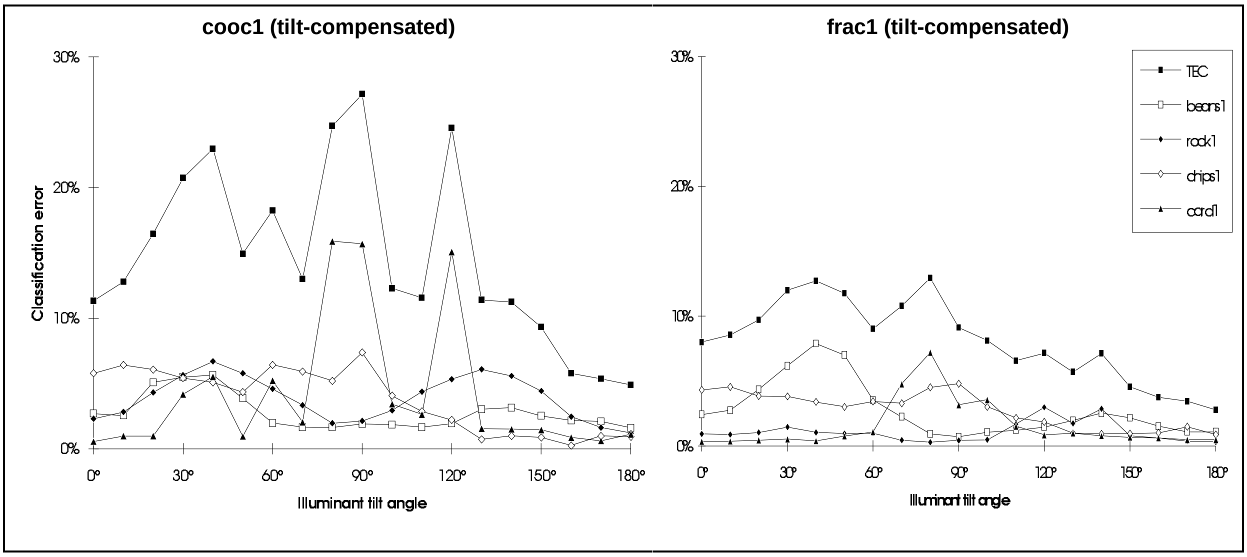

Figure 6.25 - Classification error rates of tilt-compensated cooc1 and frac1 classifiers

In contrast with the co-occurrence results the use of *F2* filters with Linnett's features (classifier *frac1*) has been more successful. Here the average error rate has been significantly reduced. Again it is the effect on the directional texture which dominates the change in the total error rates.

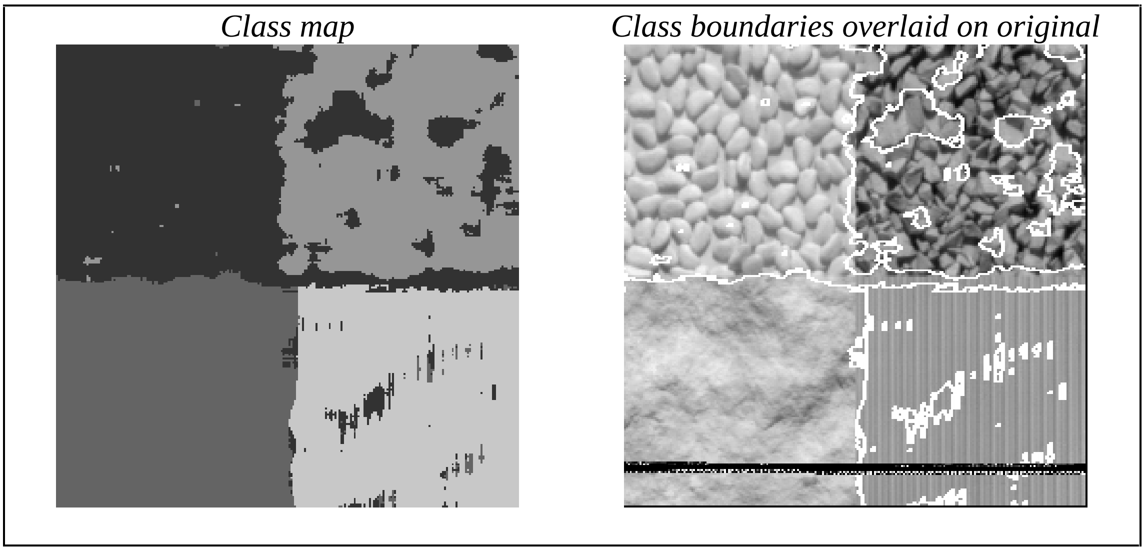

Figure 6.26 - Reduced classification error at $\tau = 90°$ (frac1 classifier)

b) A comparison of total error rates

The figures above have depicted individual error rates for each texture and total error of classification (TEC) for the three tilt-compensated classifiers. However, these graphs do not allow easy comparison of the performance of uncompensated, normalised, and tilt-compensated, classification schemes. The next three figures therefore show the total error rates that result from applying these three schemes to each classifier in turn.

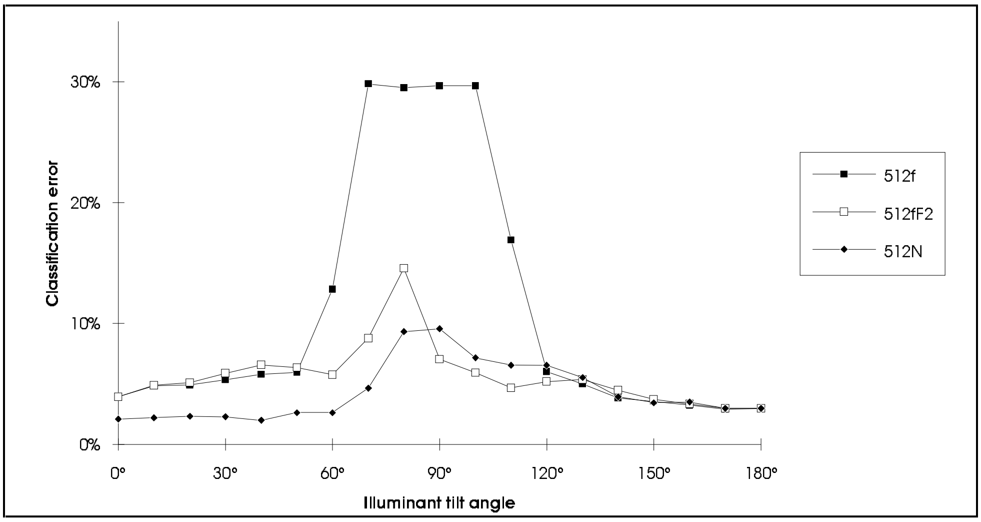

Figure 6.27 - TEC for Laws1 classifier using original(512f), normalised(512N) and tilt-compensated(512fF2) images.

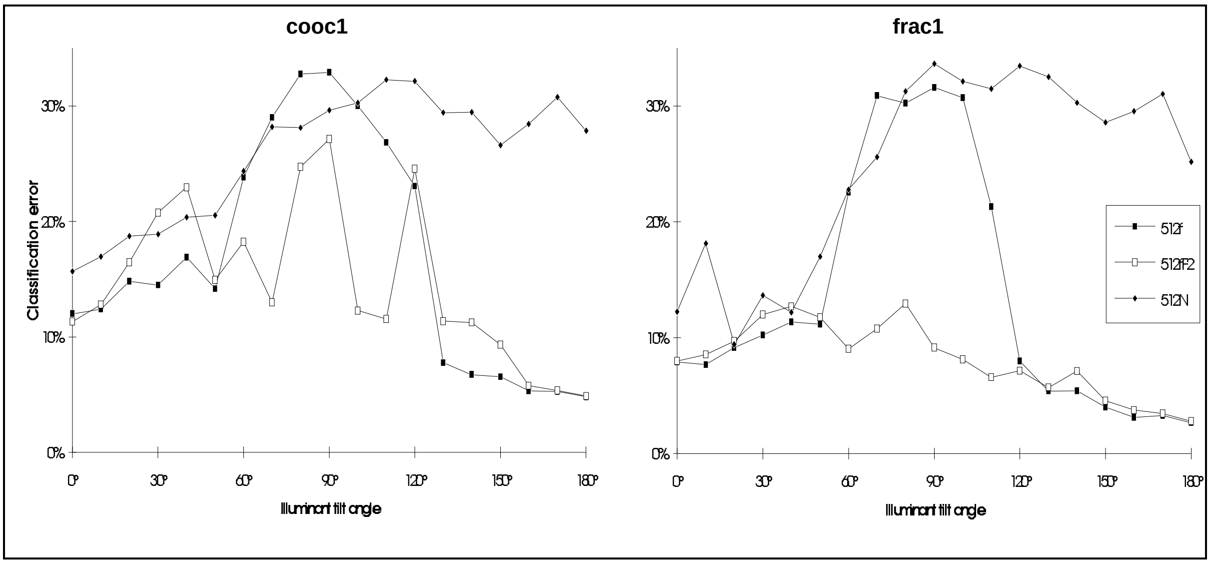

Figure 6.28 - TEC for co-occurrence (left) and Linnett's classifiers (right) using uncompensated, normalised, and tilt-compensated images (512f, 512N and 512f F2 respectively).

Figures 6.27 and 6.28 show that tilt-compensation with the *F2* filter set produces the best results with Linnett and co-occurrence features. Of the feature sets, Laws' are clearly better — with little to choose between the tilt-compensated and normalised images. However, a word of caution must be sounded — in that an alternative data set (*montage3*) showed that Laws' features with normalisation can actually increase isotropic error rates (see figure 6.11). This alternative data set was used as input to the other classifiers as well and with the exception of the normalised *Laws1* classifier results were very similar to those above. That is the tilt-compensated co-occurrence and Linnett classifiers are generally superior to their normalised counterparts.

c) Directional versus isotropic errors

What is not clear from the above graphs is the overall distribution of classification errors between directional and isotropic textures. The next figure therefore contains two graphs which show the average isotropic and directional errors for each compensation scheme. That is each point on each graph is an average calculated from the appropriate error rates of the Laws, Linnett, and co-occurrence classifiers. Thus the overall effect of each compensation scheme on isotropic and directional error rates may be examined.

From the previous theory and experimentation it might be expected that both normalisation and tilt-compensation would reduce directional texture errors, but that these two pre-processing techniques would have differing effects on isotropic textures. The graphs below, which were compiled from the results of over thirty million classification decisions, show that this is indeed the case.

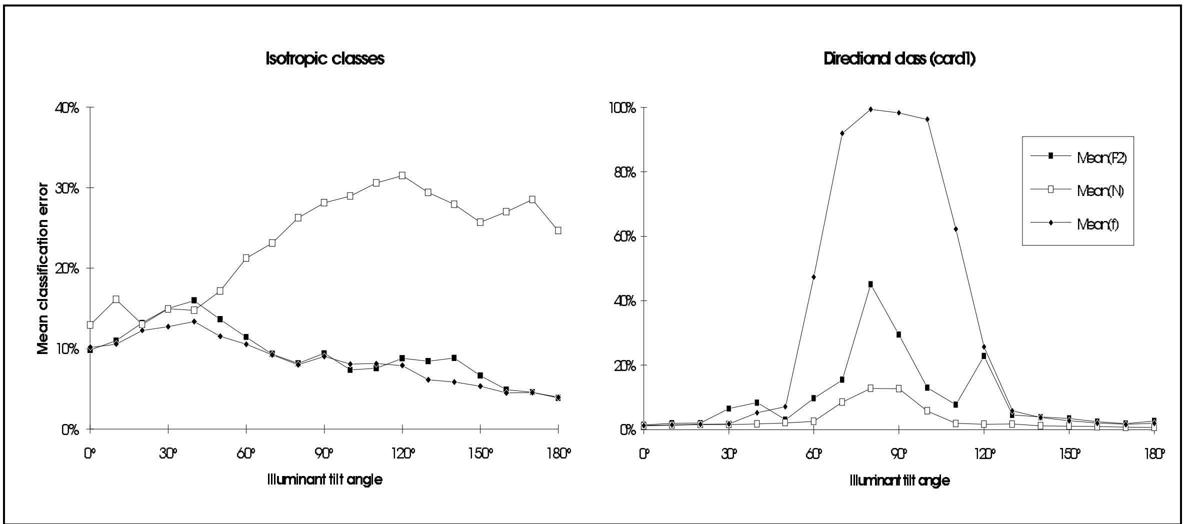

Figure 6.29 - Mean error rates for directional and isotropic texture classes (means calculated from co-occurrence, Laws', and Linnett's features; test set - montage1)

The main points that can be drawn from figure 6.29 above are :

- both normalisation and tilt-compensation reduce classification errors of the directional texture *card1,*

- normalisation produces the best results for *card1,*

- normalisation increases the average error rate of the test isotropic textures, and

- tilt-compensation does not significantly change the classifiers' ability to classify the isotropic textures of the test set.

This last point is a little disappointing given that *F2* does reduce the tilt sensitivities of all three feature sets (see table 6.3). However, the flat nature of the isotropic error rates of the uncompensated schemes (see figure 6.7) suggests that there are few tilt induced isotropic errors to compensate for in these data sets.

6.5. Conclusions

This chapter has introduced three statistical classifiers — based upon a linear discriminant and the three feature sets investigated previously. These classifiers were used to investigate the effect of variation in the direction of the illumination on supervised classification. That is the tilt and slant angles of the illuminant were varied between training and classification sessions. The test data used consisted of montages of directional and isotropic textures. The main conclusions of these investigations are as follows.

- Variation in illuminant slant between training and classification sessions induced a significant increase in the number of classification failures in all three classifiers.
- Image normalisation did not markedly reduce the number of slant induced classification errors.
- Variation in illuminant tilt between training and classification sessions significantly increased the number of mis-classifications of the directional test texture.
- Normalisation significantly reduced the number of these tilt induced errors (for the directional texture).

- However, normalisation also degraded the classifiers' performances with respect to the isotropic test textures.

In addition a tilt-compensation scheme has been developed. It is based upon an improved frequency domain model derived from four isotropic test textures. The scheme consists of a family of filters — one for each value of illuminant tilt. They are used to process images before feature generation at the training and classification stages. The conclusions drawn after testing this compensation scheme and comparing its results with those achieved with uncompensated and normalised images are :

- Tilt-compensation reduces tilt related errors for the directional texture *card1*.
- Normalisation gives better results than tilt-compensation for this directional texture.
- Unlike normalisation, tilt-compensation does not degrade a classifier's ability to classify the three isotropic textures.
- Tilt-compensation was shown to reduce the average tilt sensitivity of Laws, Linnett, and co-occurrence feature measures, when applied to the test textures.

Hence the main conclusion is that the tilt-compensation scheme developed in this chapter offers a promising method of countering variation in illuminant tilt — as it is pertinent to *both* directional *and* isotropic textures, whereas normalisation is only appropriate for directional texture.

Chapter 7

Summary and conclusions

7.1. Summary

The subjects of this thesis are : (i) the effect of variation of illuminant direction on images of topological texture, and (ii) the effect of these changes in image texture on supervised classification. In addition a tilt compensation scheme was proposed and shown to be able to reduce classification errors caused by changes in illuminant tilt between training and classification sessions.

Chapters 2 and 3 investigated (i) above. First a brief survey was presented which identified an image model of topological texture due to Kube and Pentland [Kube88]. This model was presented using a simplifying axis transformation which resulted in simpler expression for the model — providing a clearer view of its directional characteristics. It was also generalised to non-fractal surfaces. More importantly however, the implications of this model for texture classification were examined. The model predicts that the directional characteristics of image texture are not only a function of surface relief but (ignoring the more complex case uni-directional texture) are also dependent upon the tilt angle of the illumination. In addition it predicts that image variance is a function of the illuminant's slant angle. These effects are unfortunate as many texture classification schemes exploit image directionality and some exploit image variance. Variation in the latter may be removed by normalisation of images — thereby in theory removing any variations due to change in illuminant slant. Variation of image directionality due to changes in illuminant tilt may not however be compensated for in the same manner.

The third chapter used simulation and laboratory experiment to investigate the validity of this theoretical model. The results confirmed that variation in illuminant tilt can affect the directional characteristics of image texture; and that for the four isotropic test textures employed a "raised cosine" rather than the straight cosine relationship predicted by the image model was more appropriate. Future work may be able to exploit this characteristic to provide a frequency domain based illuminant tilt estimator.

The laboratory results also showed that image variance is a function of illuminant slant, but the results of simulations suggest that the predicted "sine" relationship is severely affected by shadowing. In addition the opportunity to assess the intrinsic nature of the radial shape of image magnitude spectra was taken. The gross shape of radial sections were shown to be maintained under changes in illuminant slant and tilt. However, it could not be shown that the gradient of a straight line approximating the radial section (i.e. a function of the power roll-off factor) would remain constant. Thus radial shape may, or may not, provide the basis for a practical set of texture measures that are invariant to changes in illuminant direction, but such a development must be the subject of future work.

Having established that the variance and directionality of images of isotropic texture are not invariant to changes in illuminant direction, the second part of this thesis addressed the potential impact of these variations on supervised texture classification.

Chapter 4 surveyed feature measures employed in texture classification. It identified three sets of texture measures for further investigation and in addition noted that no literature on the effects of illuminant vector variation on texture classification had been uncovered.

In chapter 5 the responses of the three sets of feature measures were investigated as regards to the effect of variation in the illuminant vector. First, the two-dimensional frequency response of the features was presented; both to provide a common view of their directionality and to give insight as to the effects predicted by the frequency domain image model presented in chapters 2 & 3. Second, this image model was used to predict the tilt response of one of Laws' feature measures — a comparison with empirical results demonstrated the image model's utility. Third, the tilt and slant responses of the features were presented. All three feature sets were affected by variation in slant. These effects were reduced by image normalisation, but this compensation was accompanied by a reduced separation between class means. The effect of tilt variation was found to depend upon the directional characteristics of the feature measure and the surface relief tested. If either were uni-directional then the output was strongly affected by illuminant tilt variation. Normalisation again reduced these effects but only if the texture was uni-directional.

Having established the effect of variation in illuminant direction on the three feature sets (as applied to the test textures) the next step was to assess the significance of these variations with respect to particular classification tasks. Hence chapter 6 introduced a simple statistical classification scheme employing a linear discriminant. This was combined with each of the three feature sets and applied to montages of isotropic and directional textures. These tests showed that variation of the illuminant's slant angle between training and classification could significantly degrade classification accuracy. Normalisation of the test sets did not markedly improve matters. Similar tests on tilt variation showed that, for the test sets employed, significant classification errors could be induced but that these were confined to the directional texture. Normalisation reduced these errors but also increased the mis-classification of isotropic textures.

The second half of chapter 6 was devoted to the development of an illuminant tilt compensation scheme. The image model developed in chapters 2 and 3 was used as the starting point from which to design a set of tilt compensation filters. In order for this scheme to be used the illuminant's tilt angle must be known during training and classification sessions — as it is used to select the appropriate filter with which to pre-process the images. Analysis of the results of employing this scheme showed that

(a) it reduced the average tilt sensitivity of all three feature sets,

(b) it reduced classification errors associated with the directional texture, and

(c) unlike normalisation it did not increase isotropic errors.

7.2. Conclusions

This thesis has used theory, simulation and laboratory experiment, to investigate the effect of changes in illuminant direction on image texture. It has been shown that variation in illuminant tilt can alter the directional properties of image texture, while variation in illuminant slant has been shown to effect a change in image variance. Both of these effects have been shown to affect the output of three sets of texture measures and hence also to significantly reduce the accuracy of classifiers employing these feature sets. To the author's knowledge the above points have not been explicitly addressed within texture classification research before.

Normalisation of images was shown to reduce variations due to changes in slant, but this was bought at the cost of reduced separation between class means of the test textures, and classification errors were not significantly reduced. Normalisation also reduced the effect of changes in tilt on images of the directional test texture. In this case classification errors associated with directional texture were reduced, whereas those associated with some of the isotropic textures increased.

A frequency domain model due to Kube and Pentland of image texture [Kube88] has, after being modified to take into account empirical observations, been used to develop a set of tilt compensation filters. Application of these filters to images of the test textures reduced the errors associated with the directional texture. Unlike normalisation it has not increased errors associated with the isotropic test textures. Neither has it reduced them. However, examination of the tilt responses suggest that there were few tilt induced isotropic errors to compensate for. In addition the reduced tilt sensitivities of tilt compensated features, when applied to isotropic textures, suggests that this scheme also has potential to improve the classification accuracy of isotropic textures.

Appendix A

The MacLaurin expansion of $(r^2+t^2+1)^{-1/2}$

The Taylor's series of a function $f(x,y)$ of two variables is

$$f(a+r,b+t)= f(a,b)+{}^{*}Df(a,b)+\frac{1}{2!}{}^{*}D^2 f(a,b)+\frac{1}{3!}{}^{*}D^3 f(a,b)...$$

Where

$${}^{*}D=\left(r\frac{\partial}{\partial x}+t\frac{\partial}{\partial y}\right)$$

Now if a = b = 0 we obtain the MacLaurin series :

$$\begin{aligned}f(r,t)&= f(0,0)+\left(r\frac{\partial}{\partial x}+t\frac{\partial}{\partial y}\right)f(0,0)\\&+\frac{1}{2!}\left(r^2\frac{\partial^2}{\partial x^2}+2rt\frac{\partial^2}{\partial x\partial y}+t^2\frac{\partial^2}{\partial y^2}\right)f(0,0)\\&+\frac{1}{3!}\left(r^3\frac{\partial^3}{\partial x^3}+3r^2t\frac{\partial^3}{\partial x^2\partial y}+3rt^2\frac{\partial^3}{\partial x\partial y^2}+t^3\frac{\partial^3}{\partial y^3}\right)f(0,0)\\&+\frac{1}{4!}\left(r^4\frac{\partial^4}{\partial x^4}+4r^3t\frac{\partial^4}{\partial x^3\partial y}+6r^2t^2\frac{\partial^4}{\partial x^2\partial y^2}+4rt^3\frac{\partial^4}{\partial x\partial y^3}+t^4\frac{\partial^4}{\partial y^4}\right)f(0,0)\\&\;\cdots\cdots\end{aligned}\qquad\text{(A.1)}$$

To find the MacLaurin expansion of

$$f(x,y)=\left(x^2+y^2+1\right)^{-\frac{1}{2}}$$

we must first find the partial derivatives at (x,y) = (0,0)

$$\begin{aligned}\frac{\partial}{\partial x}f(x,y)&=-\frac{1}{2}\left(x^2+y^2+1\right)^{-\frac{3}{2}}.2x\\&=-x\left(x^2+y^2+1\right)^{-\frac{3}{2}}\\\Rightarrow\frac{\partial}{\partial x}f(0,0)&=0\end{aligned}\qquad\text{(A.2)}$$

$$\frac{\partial^2}{\partial x^2}f(x,y)=\frac{\partial}{\partial x}\left[-x\left(x^2+y^2+1\right)^{-\frac{3}{2}}\right]$$

$$= \left\{-\left(x^2+y^2+1\right)^{-\frac{3}{2}}\right\} + \left\{-x.-\frac{3}{2}\left(x^2+y^2+1\right)^{-\frac{5}{2}}.2x\right\}$$
$$= -\left(x^2+y^2+1\right)^{-\frac{3}{2}} + 3x^2\left(x^2+y^2+1\right)^{-\frac{5}{2}}$$
$$\Rightarrow \frac{\partial^2}{\partial x^2} f(0,0) = -1 \qquad \text{(A.3)}$$

$$\frac{\partial^3}{\partial x^3} f(x,y) = \frac{\partial}{\partial x}\left[-\left(x^2+y^2+1\right)^{-\frac{3}{2}} + 3x^2\left(x^2+y^2+1\right)^{-\frac{5}{2}}\right]$$
$$= \frac{3}{2}\left(x^2+y^2+1\right)^{-\frac{5}{2}} 2x + \left\{6x\left(x^2+y^2+1\right)^{-\frac{5}{2}} + 3x^2.-\frac{5}{2}\left(x^2+y^2+1\right)^{-\frac{7}{2}}.2x\right\}$$
$$= 3x\left(x^2+y^2+1\right)^{-\frac{5}{2}} + 6x\left(x^2+y^2+1\right)^{-\frac{5}{2}} - 15x^3\left(x^2+y^2+1\right)^{-\frac{7}{2}}$$
$$= 9x\left(x^2+y^2+1\right)^{-\frac{5}{2}} - 15x^3\left(x^2+y^2+1\right)^{-\frac{7}{2}}$$
$$\Rightarrow \frac{\partial}{\partial x^3} f(0,0) = 0 \qquad \text{(A.4)}$$

$$\frac{\partial^4}{\partial x^4} f(x,y) = \frac{\partial}{\partial x}\left[9x\left(x^2+y^2+1\right)^{-\frac{5}{2}} - 15x^3\left(x^2+y^2+1\right)^{-\frac{7}{2}}\right]$$
$$= 9\left(x^2+y^2+1\right)^{-\frac{5}{2}} - 9x\frac{5}{2}\left(x^2+y^2+1\right)^{-\frac{7}{2}} 2x$$
$$-45x^2\left(x^2+y^2+1\right)^{-\frac{7}{2}} + 15x^3\frac{7}{2}.\left(x^2+y^2+1\right)^{-\frac{9}{2}}.2x$$
$$= 9\left(x^2+y^2+1\right)^{-\frac{5}{2}} - 45x^2\left(x^2+y^2+1\right)^{-\frac{7}{2}}$$
$$-45x^2\left(x^2+y^2+1\right)^{-\frac{7}{2}} + 205x^4\left(x^2+y^2+1\right)^{-\frac{9}{2}}$$
$$= 9\left(x^2+y^2+1\right)^{-\frac{5}{2}} - 90x^2\left(x^2+y^2+1\right)^{-\frac{7}{2}} - 205x^4\left(x^2+y^2+1\right)^{-\frac{9}{2}}$$
$$\Rightarrow \frac{\partial^4}{\partial x^4} f(0,0) = 9 \qquad \text{(A.5)}$$

$$\frac{\partial^2}{\partial x \partial y} f(x,y) = \frac{\partial^2}{\partial y}\left[-x\left(x^2+y^2+1\right)^{-\frac{3}{2}}\right]$$
$$= -x.-\frac{3}{2}.\left(x^2+y^2+1\right)^{-\frac{5}{2}}.2y$$
$$= 3xy\left(x^2+y^2+1\right)^{-\frac{5}{2}}$$
$$\Rightarrow \frac{\partial^2}{\partial x \partial y} f(0,0) = 0 \qquad \text{(A.6)}$$

$$\frac{\partial^3}{\partial x^2 \partial y} f(x,y) = \frac{\partial}{\partial x}\left[3xy\left(x^2+y^2+1\right)^{-\frac{5}{2}}\right]$$

$$= 3y\left(x^2 + y^2 + 1\right)^{-\frac{5}{2}} - 3xy.\frac{5}{2}\left(x^2 + y^2 + 1\right)^{-\frac{7}{2}}.2x$$

$$= 3y\left(x^2 + y^2 + 1\right)^{-\frac{5}{2}} - 15x^2y\left(x^2 + y^2 + 1\right)^{-\frac{7}{2}}$$

$$\Rightarrow \frac{\partial^3}{\partial x^2 \partial y} f(0,0) = 0 \tag{A.7}$$

$$\frac{\partial^4}{\partial x^3 \partial y} f(x,y) = \frac{\partial}{\partial y}\left[9x\left(x^2 + y^2 + 1\right)^{-\frac{5}{2}} - 15x^3\left(x^2 + y^2 + 1\right)^{-\frac{7}{2}}\right]$$

$$= \left[9x.-\frac{5}{2}.\left(x^2 + y^2 + 1\right)^{-\frac{7}{2}}.2y\right] - \left[15x^3.-\frac{7}{2}\left(x^2 + y^2 + 1\right)^{-\frac{9}{2}}.2y\right]$$

$$= -45xy\left(x^2 + y^2 + 1\right)^{-\frac{7}{2}} + 105x^3y\left(x^2 + y^2 + 1\right)^{-\frac{9}{2}}$$

$$\Rightarrow \frac{\partial^3}{\partial x^2 \partial y} f(0,0) = 0 \tag{A.8}$$

$$\frac{\partial^4}{\partial x^2 \partial y^2} f(x,y) = \frac{\partial}{\partial y}\left[3y\left(x^2 + y^2 + 1\right)^{-\frac{5}{2}} - 15x^2y\left(x^2 + y^2 + 1\right)^{-\frac{7}{2}}\right]$$

$$= 3(x^2 + y^2 + 1)^{-\frac{5}{2}} - 15y^2(x^2 + y^2 + 1)^{-\frac{7}{2}}$$

$$- 15x^2(x^2 + y^2 + 1)^{-\frac{7}{2}} - 105x^2y^2(x^2 + y^2 + 1)^{-\frac{9}{2}}$$

$$\Rightarrow \frac{\partial^4}{\partial x^2 \partial y^2} f(0,0) = 3 \tag{A.9}$$

Substituting the partial derivatives (A.2) to (A.9) in the MacLaurin expansion (A.1) gives:

$$f(r,t) = 1 + [r.0 + t.0] + \frac{1}{2!}\left[-r^2 + 2rt.0 - t^2\right]$$

$$+ \frac{1}{3!}\left[r^3.0 + 3r^2t.0 + 3rt^2.0 + t^3.0\right]$$

$$+ \frac{1}{4!}\left[r^4.9 + 4r^3t.0 + 6r^2t^2.0 + 4rt^3.0 + t^4.9\right]$$

$$+ \cdots\cdots$$

$$\Rightarrow \left(r^2 + t^2 + 1\right)^{-1/2} = 1 - \frac{1}{2!}\left(r^2 + t^2\right) + \frac{9}{4!}\left(r^2 + t^2\right)^2 \cdots\cdots$$

References

Ahuja81 N. Ahuja & B.J. Schachter, "Image models", Computing Surveys, V13, December 1981, pp373-397.

Arduini92 F. Arduini, S. Fioravanti, D.D. Giusto, & F. Inzirillo, "Multifractals and texture classification", Proceedings of IEE International Conference on Image Processing and its Applications, Maastricht, April 1992, pp454-457.

Bajcsy76 R. Bajcsy & L. Lieberman, "Texture gradient as a depth cue", CGIP, V5, 1976, pp52-67.

Baston75 R.M. Baston, K. Edwards & E.M. Edwards & E.M Eliason, "Computer generated shaded-relief images", Journal U.S. Geol. Survey, V3 Pt 4, July/August 1975, pp401-408 .

Beckmann63 P. Beckmann & A. Spizzichino, "*The scattering of electromagnetic waves from rough surfaces*", Pergamon, 1963.

Bennett89 J.M. Bennett & L. Mattsson, "*Surface roughness and scattering*", Optical Society of America, Washington D.C., 1989.

Blake90 A. Blake & C. Marinos, "Shape from texture : estimation, isotropy and moments", Artificial Intelligence, V45, 1990, pp323-380.

Blinn78 J. Blinn, "Simulation of wrinkled surfaces", Proceedings of SIGRAPH '78, 1978, pp286-292.

Blinn90 J. Blinn, "The truth about texture mapping", IEEE Computer Graphics & Applications, March 1990, pp78-83.

Blostein89 D. Blostein & N. Ahuja, "Shape from texture: integrating texture-element extraction and surface estimation", PAMI, V11, December 1989, pp1233-1251.

Boulanger89 P. Boulanger, A. Gagalowicz & M. Rioux, "Integration of synthetic surface relief in range images", CVGIP V47, 1989, pp361-372.

Bouville85 C. Bouville, "Bounding ellipsoids for ray-fractal intersection", SIGGRAPH, V19, July 1985, pp45-52.

Bovik87 A.C. Bovik, M. Clark & W.S. Geisler, "Computational texture analysis using localised spatial filtering", Proceedings of IEEE Computer Society workshop on Computer Vision, Miami Beach, 1987, pp201-206.

Box76 G.E.P. Box & G.M. Jenkins "*Time series analysis forecasting and control*", Holden-Day, San Francisco, 1976.

Brigham88 E.O. Brigham, "*The fast Fourier transform and its applications*", Prentice Hall, 1988.

Brodatz66 P. Brodatz, "*Textures : a photographic album for artists and designers*", Dover, New York, 1966.

Castrec88 P. Castrec & J-P. Kernin, "Textural segmentation of natural sea bottom", Proceedings of Undersea Defence Technology Conference, Paris, 1988, pp445-450.

Chantler91 M.J. Chantler, "Texture analysis of underwater video images", Research memorandum RM/91/25, Dept. of Electrical & Electronic Engineering, Heriot-Watt University, December 1991.

Chellappa82 R. Chellappa & R.L. Kashyap, "Digital image restoration using spatial interaction models", ASSP V30, June 1982, pp461-472.

Chellappa85a R. Chellappa & S. Chatterjee, "Classification of textures using Gaussian Markov random fields", ASSP V33, August 1985, pp959-963.

Chellappa85b R. Chellappa, S. Chatterjee & R. Bagdazian, "Texture synthesis and compression using Gaussian-Markov random fields", SMC, V15, No. 2, 1985, pp298-308.

Chen90 S.S. Chen, J.M. Keller & R.M. Crownover, "Shape from fractal geometry", Artificial Intelligence, V43, 1990, pp199-218.

Choe91a Y. Choe & R.L. Kashyap, "Modeling, estimation, and pattern analysis of random texture on 3-D surfaces", Report TR-EE-91-4, School of Electrical Engineering, Purdue University, Indiana, January 1991.

Choe91b Y. Choe & R.L. Kashyap, "3-D shape from a shaded and textural surface image", PAMI V13, Sept. 1991, pp907-919.

Clarke92 S.J. Clarke, "The analysis and synthesis of texture in sidescan sonar data", Ph.D. thesis, Dept. Electrical & Electronic Engineering, Heriot-Watt University, May 1992.

Cohen91a F.S. Cohen, Z. Fan & M.A.S. Patel, "Classification of rotated and scaled textured images using Gaussian Markov field models", PAMI V13, February 1991, pp192-202.

Cohen91b F.S. Cohen, Z. Fan & S. Attali, "Automated inspection of textile fabrics using textural models", PAMI V13, August 1991, pp803-808.

Cohen91c F.S. Cohen & M.A.S. Patel, "Modeling and synthesis of images of 3d textured surfaces", CVGIP: graphical models and image processing, V53, November 1991, pp501-510.

Conners80 R.W. Conners & C.A. Harlow, "A theoretical comparison of texture algorithms", PAMI V2, May 1980, pp204-222.

Cooper86 G.R. Cooper & C.D. McGillem, "*Probabilistic methods of signal and system analysis*", CBS College Publishing, 1986.

D'Astous84 F. D'Astous and M.E. Jernigan, "Texture discrimination based on detailed measures of the power spectrum", 7th International Conference on Pattern Recognition, Montreal, July 1984, pp83-86.

Davies90 E.R. Davies, "*Machine vision: theory, algorithms, practicalities*", Academic Press, 1990.

Davis73 J.C. Davis, "*Statistics and data analysis in Geology*", Wiley, 1973.

Davis81a L.S. Davis, M. Clearman & J.K. Aggarwal, "An empirical evaluation of generalised cooccurrence matrices", PAMI V3, March 1981, pp214-221.

Davis81b L.S. Davis, "Polarograms : a new tool for image texture analysis", Pattern Recognition, V13, No. 3, 1981, pp219-223.

Davis83 L.S. Davis, L. Janos & S.M. Dunn, "Efficient recovery of shape from texture", PAMI V5, Sept. 1983, pp485-492.

Dennis89 T.J. Dennis & N.G. Dessipris, "Fractal modelling in image texture analysis", IEE Proceedings, V136, Pt. F, No. 5, October 1989, pp227-235.

Dinstein84 I. Dinstein, A.C. Fong, L.M. Ni & K.Y. Wong, "Fast discrimination between homogeneous and textured regions", Proceedings of IEEE 7th International Conference on Pattern Recognition , Montreal, July 1984, pp361-363.

duBuf90 J.M.H. du Buf, M. Kardan & M. Spann, "Texture feature performance for image segmentation", Pattern Recognition, V23, No. 3/4, 1990, pp291-309

Gagalowicz86 A. Gagalowicz & S. De Ma, "Model driven synthesis of natural textures for 3-D scenes", Computers & Graphics, V10, 1986, pp161-170.

Gibson50 J. Gibson, "*The perception of the visual world*", Houghton Mifflin Company, Boston, MA, 1950.

Greenhill93 D. Greenhill & E.R. Davies, "Texture analysis using neural networks and mode filters", Proceedings of BMVC93, September 1993, University of Surrey, Guildford, pp509-518.

Haralick73 R.M. Haralick, K. Shanmugam & I. Dinstein, "Textural features for image classification", SMC, November 1973, pp611-621.

Haralick79 R.M. Haralick, "Statistical and structural approaches to texture", Proceedings of the IEEE, V67, No. 5, May 1979, pp786-804.

Haruyama84 S. Haruyama & B.A. Barsky, "Using stochastic modeling for texture generation", IEEE Computer Graphics & Applications, March 1984, pp7-19.

Harwood85 D. Harwood, M. Subbarao & L.S. Davis, "Texture classification by local rank correlation", CVGIP V32, 1985, pp404-411.

Hassner80 M. Hassner & J. Sklansky, "The use of Markov random fields as models of texture", CGIP V12, 1980, pp357-370.

He88 D-C He, L. Wang & J. Guibert, "Texture discrimination based on an optimal utilization of texture features", Pattern Recognition, V21, No. 2, 1988, pp141-146.

Heckbert86 P.S. Heckbert, "Survey of texture mapping", IEEE Computer Graphics & Applications, November 1986, pp56-67.

Horn89 B.K.P. Horn & M.J. Brooks, "*Shape from shading*", MIT Press, 1989.

Hosking81 J.R.M. Hosking, "Fractional differencing", Biometrika, V68, 1981, pp165-176.

Huang72 T.S. Huang, "Two-dimensional windows", IEEE Transactions on Audio and Electroacoustics, March 1972, pp88-89.

Huang78 T.S. Huang, G.J. Yang & G.Y. Tang, "A fast two-dimensional median filtering algorithm", Proceedings of IEEE Conference on Pattern Recognition and Image Processing, May, 1978, pp128-131.

Ikeuchi81 K. Ikeuchi & B.K.P. Horn, "Numerical shape from shading and occluding boundaries", Artificial Intelligence, V 17, p141-184.

Jackson86 D.R. Jackson, D.P. Winebrenner & A. Ishimaru, "Application of the composite roughness model to high-frequency bottom backscattering", Journal Acoustic Society of America, V79, No. 5, May 1986, pp1410-1422.

Jain91 A.K. Jain & F. Farrokhnia, "Unsupervised texture segmentation using Gabor filters", Pattern Recognition, V24, No. 12, 1991, pp1167-1186.

James85 M. James, "*Classification algorithms*", Collins, 1985.

Kanatani84 K. Kanatani, "Detection of surface orientation and motion from texture by a stereological technique", Artificial Intelligence, V23, 1984, pp213-237.

Kartikeyan91 B. Kartikeyan & A. Sarkar "An identification approach for 2-d autoregressive models in describing textures", CVGIP Graphical Models and Image Processing, V53, March 1991, pp121-131.

Kashyap80 R.L. Kashyap, "Univariate and multivariate random field models for images", CGIP V12, 1980, pp257-270.

Kashyap83 R.L. Kashyap & R. Chellappa, "Estimation and choice of neighbors in spatial-interaction models of images", IT V29, January 1983, pp60-72.

Kashyap84 R.L. Kashyap & P.M. Lapsa, "Synthesis and estimation of random fields using long correlation models", PAMI V6, November 1984, pp800-809.

Kashyap86 R.L. Kashyap & A. Khotanzad, "A model-based method for rotation invariant texture classification", PAMI V8, July 1986, pp472-481.

Kashyap89 R.L. Kashyap & K-B Eom, "Texture boundary detection based on the long correlation model", PAMI V11, January 1989, pp58-67.

Kay81 S.M. Kay & S.L. Marple, "Spectrum analysis - a modern perspective", Proceedings IEEE, V69, No. 11, November 1981, pp1380-1419.

Keller87 J.M. Keller, R.M. Crownover, & R.Y. Chen, "Characteristics of natural scenes related to the fractal dimension" PAMI, V9, Sept.1987, pp621-627.

Keller89 J.M. Keller, & S. Chen, "Texture description and segmentation through fractal geometry" CVGIP V45, 1989, pp150-166.

Kender80 J.R. Kender, "Shape from Texture", Ph.D. thesis, Carnegie-Mellon University, November 1980.

Khotanzad87 A. Khotanzad and R.L. Kashyap, "Feature Selection for Texture Recognition Based on Image Synthesis", SMC V17, November/December 1987, pp1087-1095.

Khotanzad89 A. Khotanzad and J-Y Chen, "Unsupervised Segmentation of Textured Images by Edge Detection in Multidimensional Features", PAMI V11, April 1989, pp414-421.

Kruger74 R.P. Kruger, W.B. Thompson & A.F. Turner, "Computer diagnosis of Pneumoconiosis" SMC V4, 1974, pp40-49.

Kube88 Kube, P. & Pentland, A., "On the imaging of fractal surfaces", IEEE Transactions Pattern Analysis and Machine Intelligence, Vol. 10, No. 5, pp704-707, Sept. 1988.

Laws79 K.I. Laws, "Texture energy measures", Proceedings of Image Understanding Workshop, November 1979, pp47-51.

Laws80 K.I. Laws, "Textured image segmentation", Ph.D. thesis, Dept. Electrical Engineering, University of Southern California, January 1980.

Lendaris70 G.G. Lendaris & G.L. Stanley, "Diffraction-pattern sampling for automatic pattern recognition" Proceedings of IEEE, V58, February 1970, pp198-216.

Lim90 J.S. Lim, "*Two-dimensional signal and image processing*", Prentice-Hall, 1990.

Linnett91a L.M. Linnett, "Multi-texture image segmentation", Ph.D. thesis, Dept. Computing and Electrical Engineering, Heriot-Watt University, March 1991.

Linnett91b L.M. Linnett, S.J. Clarke, C. Graham & D.N. Langhorne, "Remote sensing of the sea-bed using fractal techniques", Electronics & Communication Engineering Journal, October 1991, pp195-203.

Linnett93 L.M. Linnett, D.R. Carmichael, S.J. Clarke & A.D. Tress, "Texture analysis of sidescan sonar data", IEE seminar "Texture analysis in radar and sonar", London, November 1993, pp2/1-2/6.

Lovell92	R. Lovell, W.R. Uttal, T. Shepherd & S. Dayanand, "A model of visual texture discrimination using multiple weak operators and spatial averaging", Pattern Recognition, V25, No. 10, 1992, pp1157-1170.
McDaniel83	S.T. McDaniel & A.D. Gorman, "An examination of the composite-roughness scattering model", Journal Acoustic Society of America, V73, No. 5, May 1983, pp1476-1486.
Mandelbrot83	B.B. Mandelbrot, "*The fractal geometry of nature*", W.H. Freeman, New York, 1983.
Mandelbrot85	Mandelbrot, B.B., "Self-affine fractals and the fractal dimension", Physica Scripta V32, pp257-260, 1985.
Mao92	J. Mao & A.K. Jain, "Texture classification and segmentation using multiresolution simultaneous autoregressive models", Pattern recognition, V25, No.2, 1992, pp 173-188.
Marple87	S.L. Marple, "*Digital spectral analysis with applications*", Prentice Hall, 1987.
Medioni84	G.G. Medioni & Y. Yasumoto, "A note on using the fractal dimension for segmentation", Proceedings of the workshop on Computer Vision: Representation and Control, Annapolis, MD, USA, April-May 1984, pp25-30.
Miller91	P. Miller & S. Astley, "Classification of breast tissue by texture analysis", Proceedings of BVMC91, Glasgow, Sept. 1991.
Mosquera92	A. Mosquera, D. Cabello, M.J. Carreira & M.G. Penedo, "A fractal-based approach to segmentation", Proceedings of IEE International Conference on Image Processing and its Applications, Maastricht, Netherlands, April 1992, pp450-453.
Newman79	W.M. Newman & R.F. Sproull, "*Principles of interactive computer graphics*", McGraw Hill, 1979.
Ogilvy87	J.A. Ogilvy, "Wave scattering from rough surfaces", Reports on Progress in Physics, V50, 1987, pp1553-1608.
Ogilvy91	J.A. Ogilvy, "*Theory of wave scattering from random rough surfaces*", Adam Hilger, 1991.
Patel91	Patel, A.S., "Shape extraction of 3-D textured scenes from image data", Ph.D. thesis, Drexel University, July 1991.
Patel93	M.A.S. Patel & F.S. Cohen, "Local surface shape estimation of 3-D textured surfaces using Gaussian Markov Random Fields and stereo windows", PAMI, V15, No. 10, October 1993, pp 1091-1098.
Peebles87	P.Z. Peebles, "*Probability, random variables, and random signal principles*", McGraw-Hill, 1987.
Peleg84	S. Peleg, J. Naor, R. Hartley & D. Avnir, "Multiple resolution texture analysis and classification", PAMI V6, July 1984, pp518-523.
Peli90	T. Peli, "Multiscale fractal theory and object characterisation", Journal Optical Society of America A, V7, No. 6, June 1990, pp1101-1112.
Pentland84	A.P. Pentland, "Fractal based description of natural scenes", PAMI V6, 1984, pp661-674.
Pentland86	A.P. Pentland, "Shading into texture", Artificial Intelligence, V29, 1986, pp147-170.

Pietikainen82 M. Pietikainen, A. Rosenfeld & L.S. Davis, "Texture classification using averages of local pattern matches", Proceedings of IEEE Computer. Society Conference on Pattern Recognition and Image Processing, 1982.

Pietikainen83 M. Pietikainen, A. Rosenfeld & L.S. Davis, "Experiments with texture classification using averages of local pattern matches", SMC V13 May/June 1983, pp421-426.

Reed93 T.R. Reed & J.M. Hans du Buf, "A review of recent texture segmentation and feature extraction techniques" CVGIP: Image Understanding, V57, No. 3, pp359-372, May 1993.

Rogers85 D.F. Rogers, "*Procedural elements for computer graphics*", McGraw Hill, 1985.

Rosenfeld75 A. Rosenfeld, "A note on the automatic detection of texture gradients", IEEE Transactions on Computers, October 1975, pp988-991.

Saupe88 D. Saupe, "Algorithms for random fractals" in *The science of fractal images*, H-O Peitgen & D Saupe (eds.), Springer Verlag, 1988, pp71-113.

Shang93 C. Shang & K. Brown, "Texture classification of side-scan sonar images with neural networks", IEE seminar "Texture analysis in radar and sonar", London, November 93, pages 3/1 to 3/9.

Sun83 C. Sun & W.G. Wee, "Neighboring gray level dependence matrix for texture classification", CVGIP V23, 1983, pp341-352.

Tou74 J.T. Tou & R.C. Gonzalez, "*Pattern recognition principles*", Addison-Wesley, 1974.

Unser86 M. Unser, "Sum and difference histograms for texture classification" PAMI V8, January 1986, pp118-125.

VanGool85 L. Van Gool, P. Dewaele & A. Oosterlinck, "Texture analysis anno 1983", CVGIP V29, 1985, pp336-357.

Vorburger93 T.V. Vorburger, E. Marx & T.R. Lettieri, "Regimes of surface roughness measurable with light scattering", Applied Optics, V32, No. 19, July 1993, pp3401-3408.

Voss88 R.F. Voss, "Fractals in nature: from characterisation to simulation" in *The science of fractal images*, H-O. Peitgen & D. Saupe (eds.), Springer Verlag, 1988, pp21-70.

Wechsler80 H. Wechsler, "Texture analysis - a survey" Signal Processing 2, pp271-282, 1980.

Welch67 P.D. Welch, "The use of fast Fourier transform for the estimation of power spectra : a method based on time averaging over short, modified periodograms", IEEE Transactions Audio and Electroacoustics, AU-15, No. 2, June 1967, pp70-73.

Weszka76 J.S. Weszka, C.R. Dyer 7 A. Rosenfeld, " A comparative Study of Texture Measures for Terrain Classification" SMC, April 1976, pp269-285.

Witkin81 A.P. Witkin, "Recovering surface shape and orientation from texture", Artificial Intelligence, V17, 1981, pp17-45.

Zucker80 S.W. Zucker & D. Terzopoulos, "Finding Structure in Co-occurrence Matrices for Texture Analysis", CGIP, V12, 1980, pp286-308.

Abbreviations for references

ASSP	-	IEEE Transactions on Acoustics, Speech, and Signal Processing
BMVC	-	British Machine Vision Conference
CGIP	-	Computer, Graphics & Image Processing
CVGIP	-	Computer Vision, Graphics & Image Processing
IT	-	IEEE Transactions on Information Theory
PAMI	-	IEEE Transactions on Pattern Analysis and Machine Intelligence
SMC	-	IEEE Transactions on Systems, Man, & Cybernetics