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 :

$$H(\omega_1,\omega_2) = H_1(\omega_1)H_2(\omega_2) \qquad (5.1)$$

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 :

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

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

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

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

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

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).

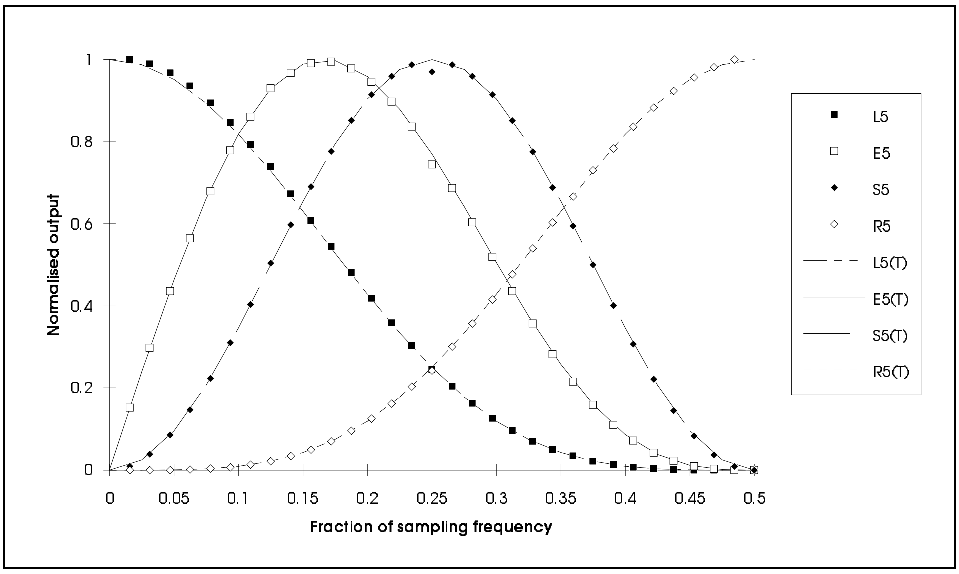

Figure 5.2 - Laws' one-dimensional operators : observed and theoretical (T) frequency responses

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.

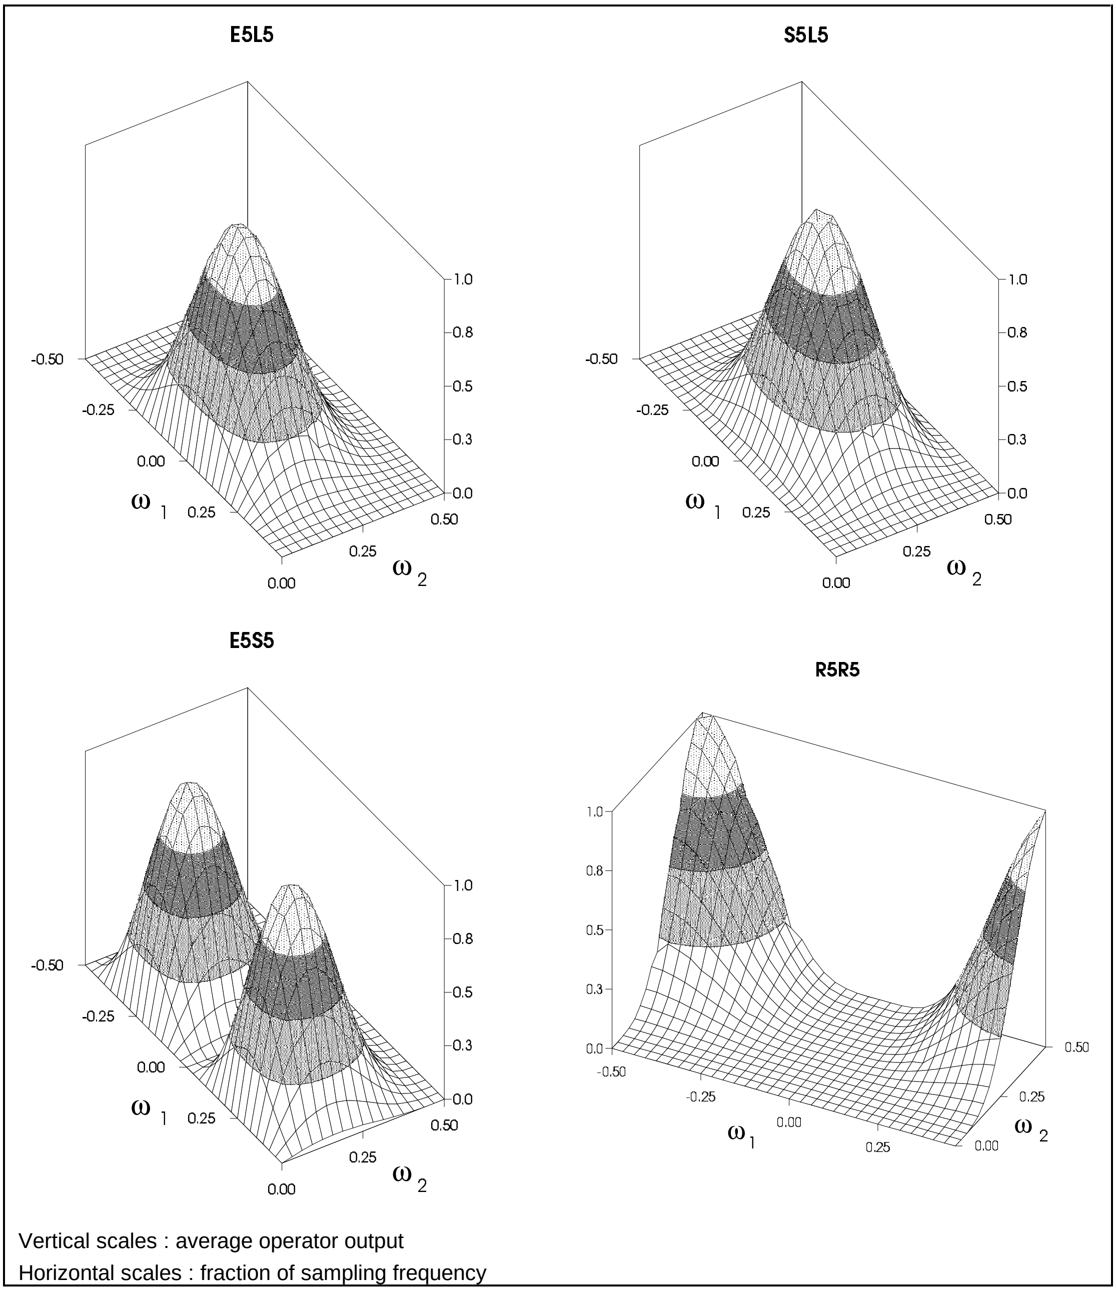

Figure 5.3 - Laws' operators : empirical two-dimensional frequency responses

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 :

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

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 :

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

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

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

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_1) 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.

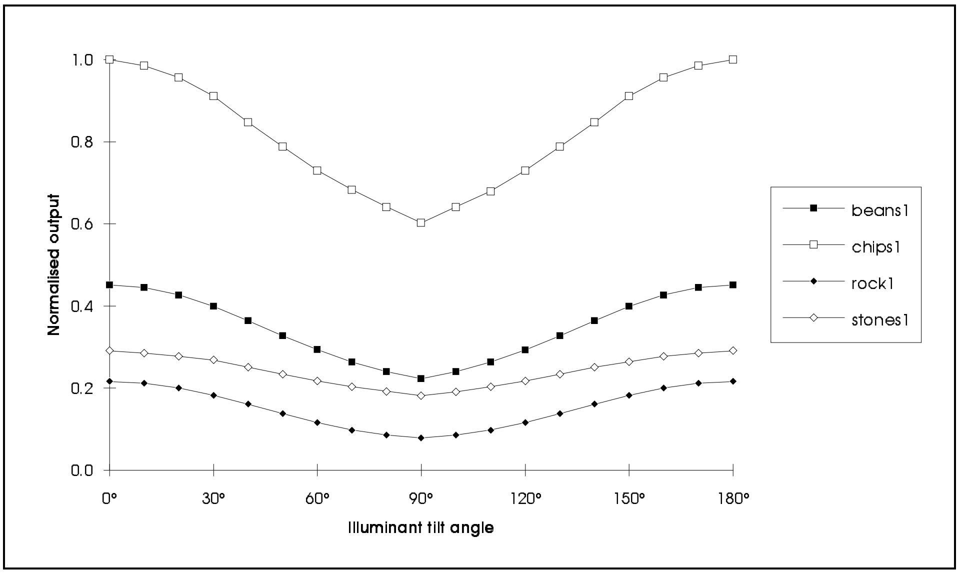

Figure 5.4 - Predicted effect of tilt angle variation on L5E5 output[2]

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.

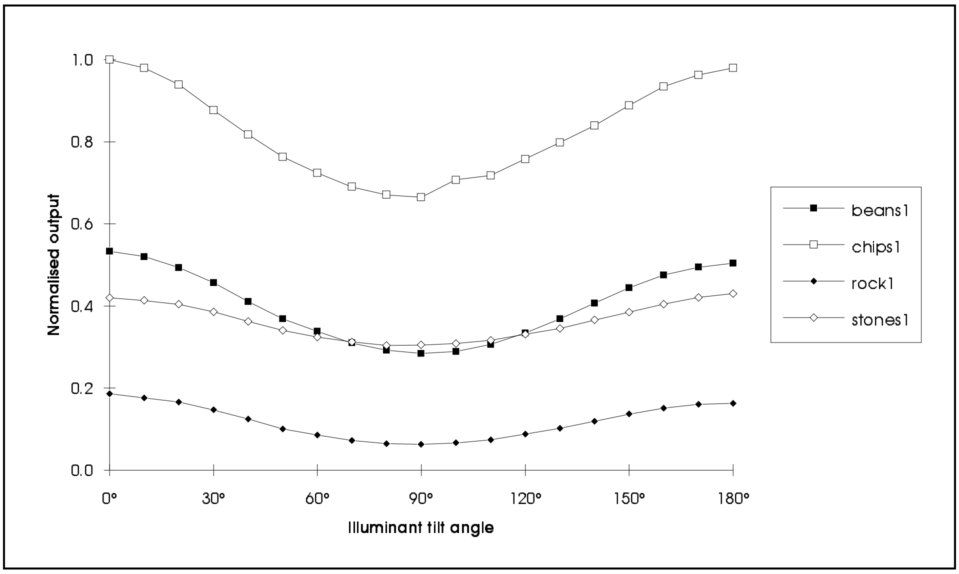

Figure 5.5 - Observed effect of tilt angle variation on L5E5 output[2]

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.

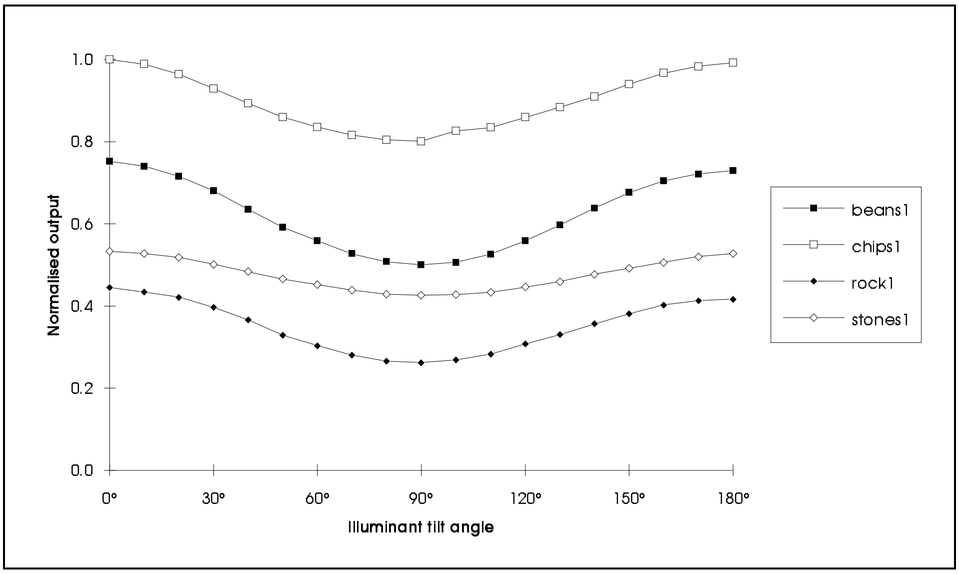

Figure 5.6 - Effect of tilt angle variation on L5E5 output (ABSAVE macro statistic)[2]

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.

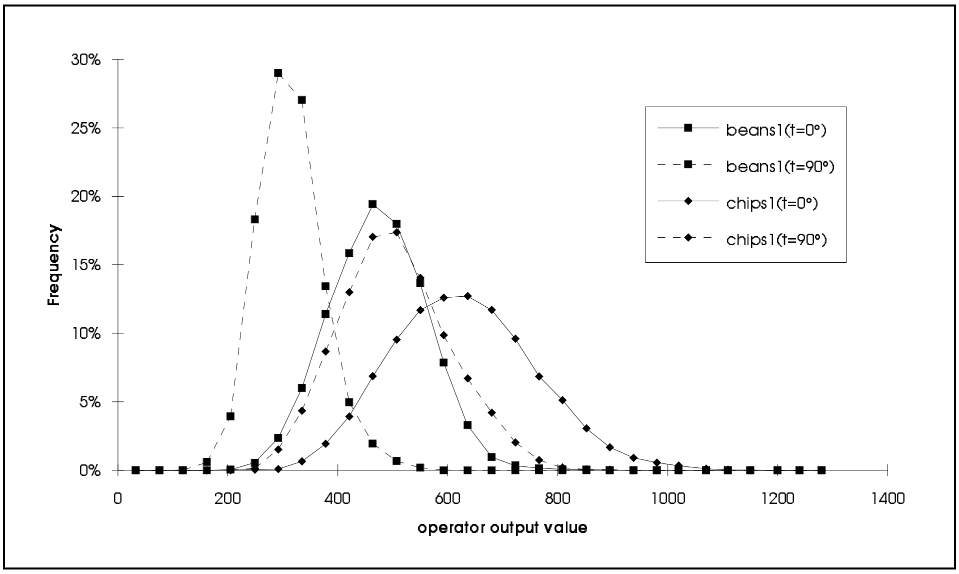

Figure 5.7 - Effect of tilt variation on L5E5 distributions

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*).

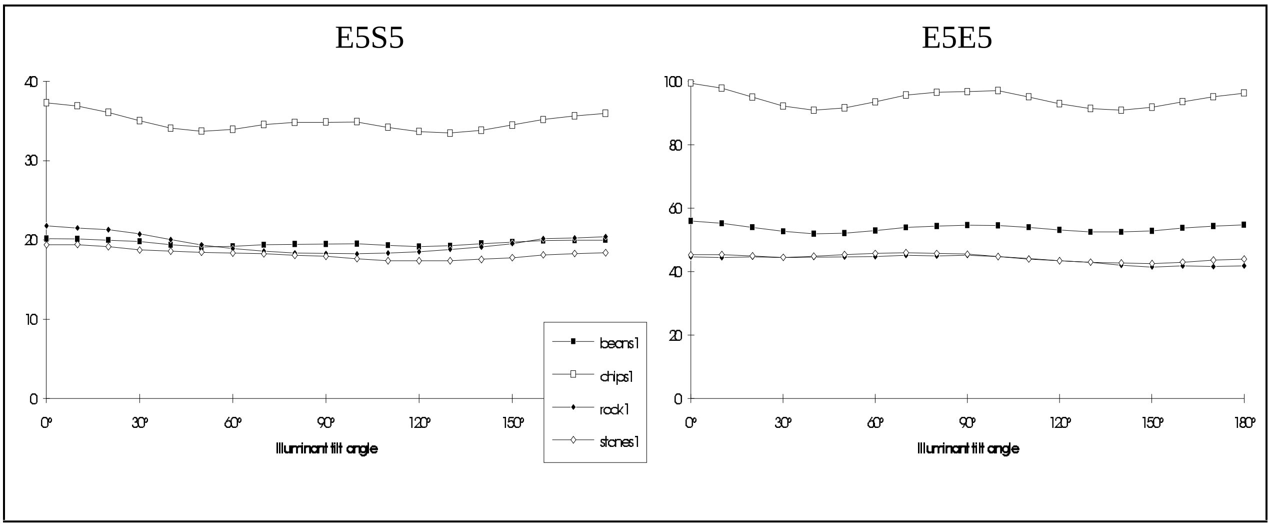

Figure 5.8 - Tilt response of the bi-directional operators E5S5 and E5E5

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.

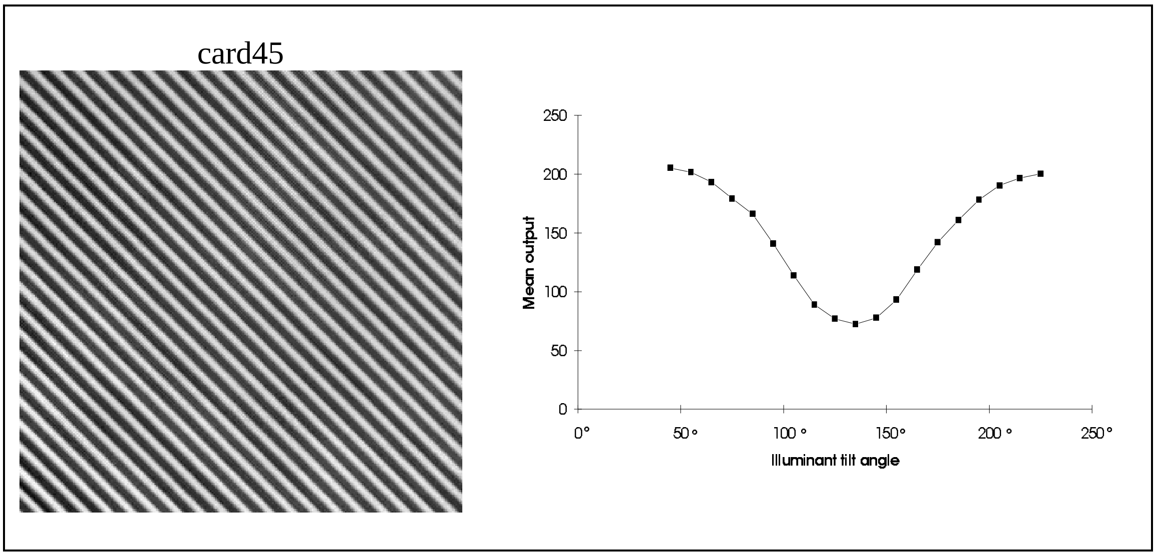

Figure 5.9 - The tilt response of E5E5 for the directional texture "card45"

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.

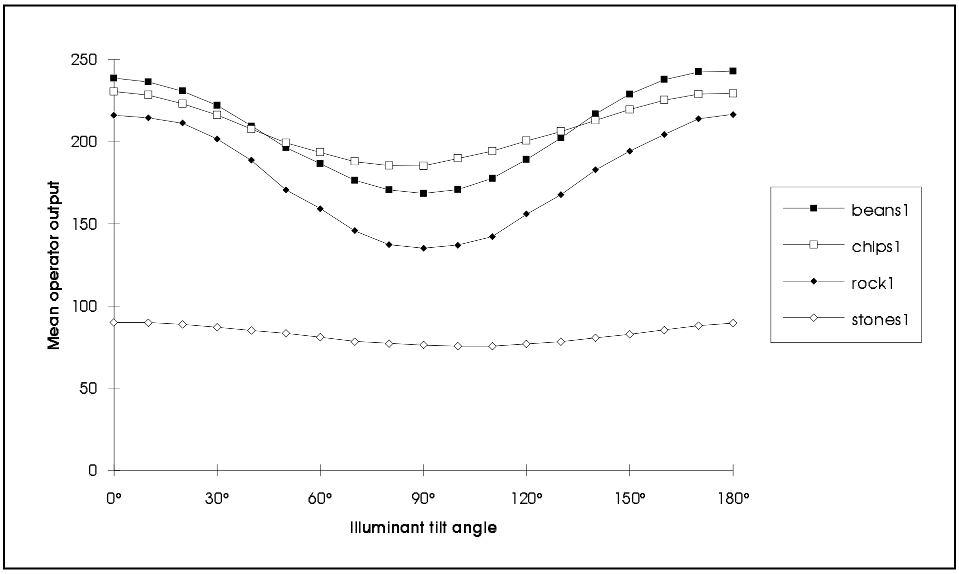

Figure 5.10 - The effect of normalisation on L5E5 tilt angle response

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.

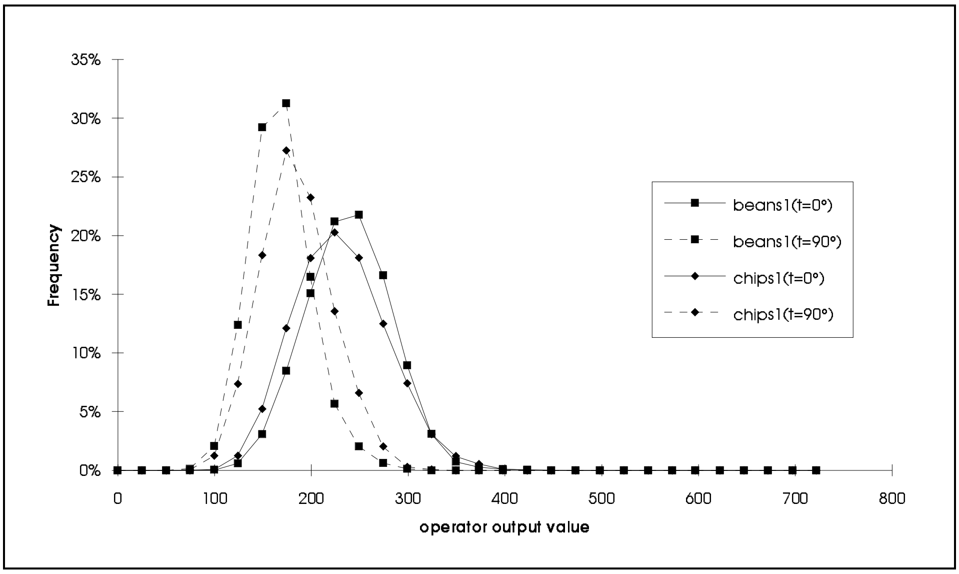

Figure 5.11 - Effect of normalisation on the L5E5 output distribution for tilts of 0° and 90°.

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.

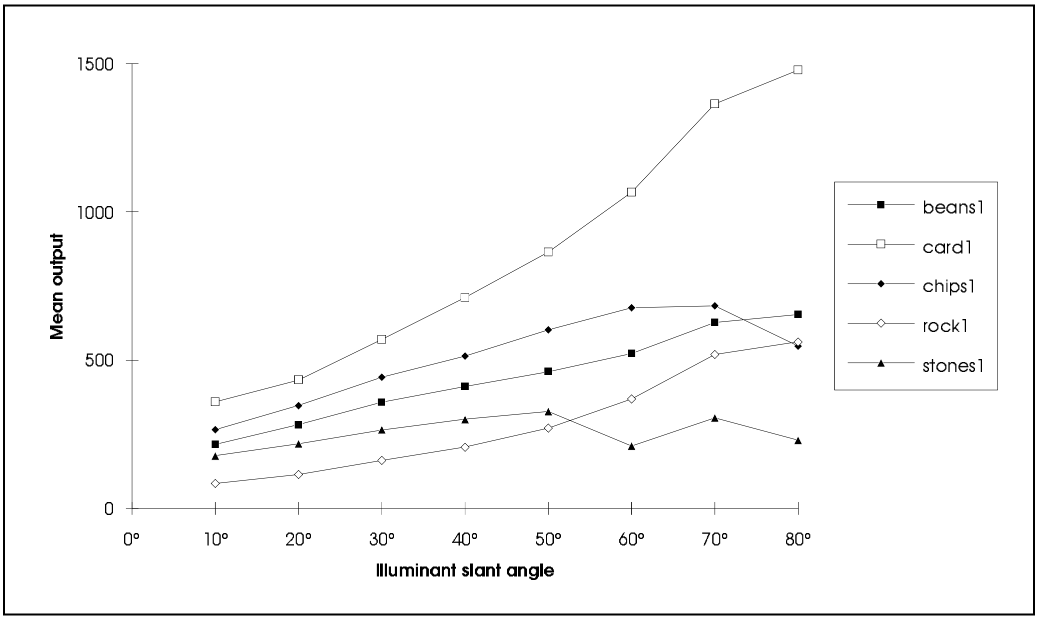

Figure 5.12 - Un-normalised slant angle response of Laws' L5E5 operator

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.

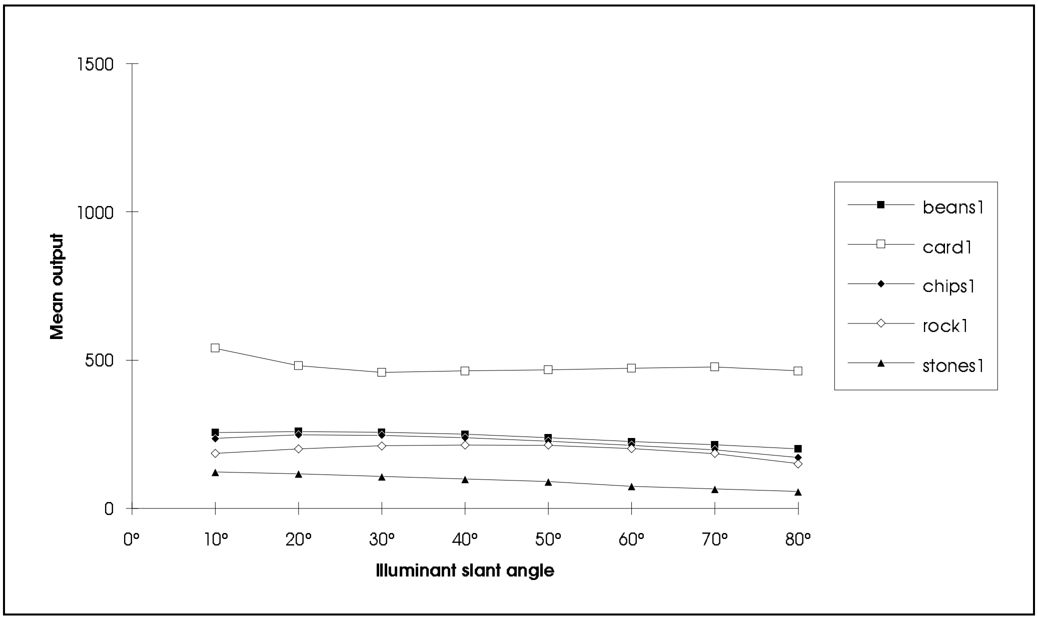

Figure 5.13 - Normalised response of Laws' L5E5 operator

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.

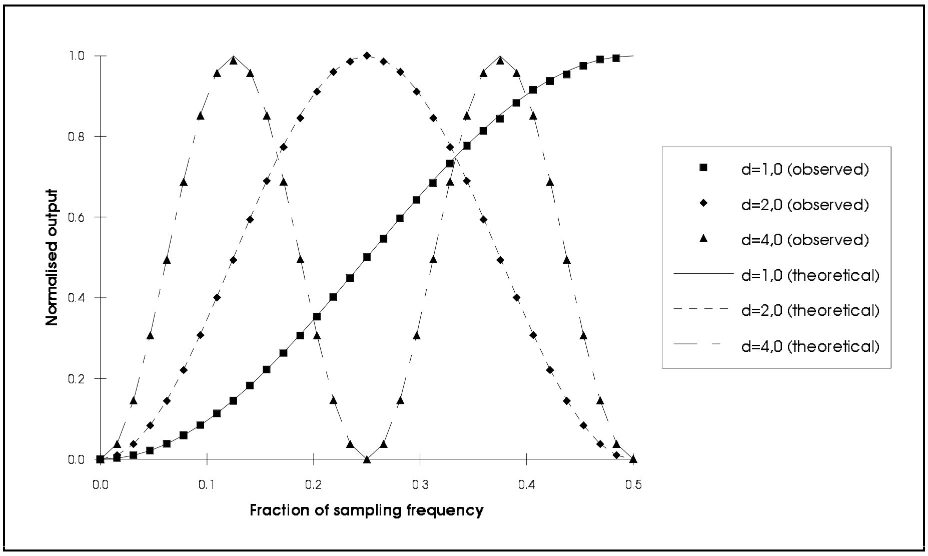

Figure 5.14 - CON operator : effect of changing the size of the displacement vector

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.

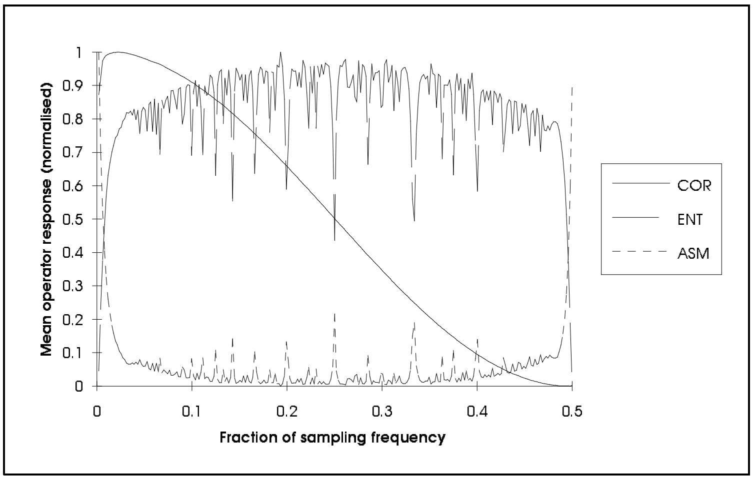

Figure 5.15 - One-dimensional frequency responses of co-occurrence features (d=(1,0), Ng = 16)

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.

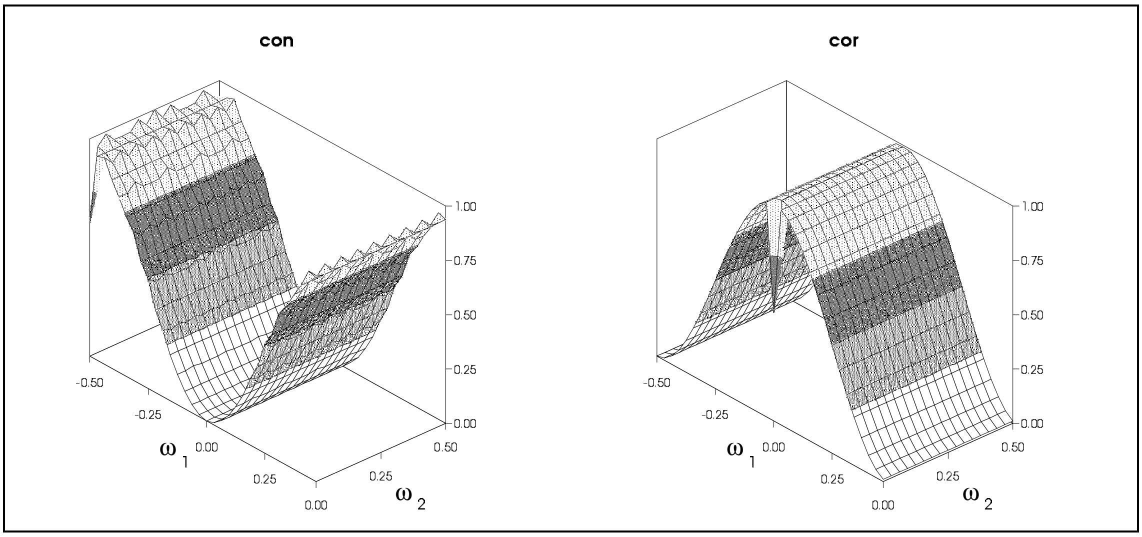

Figure 5.16 - Two-dimensional frequency responses of the co-occurrence contrast and correlation features for a displacement vector **d** = (0,1)

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.

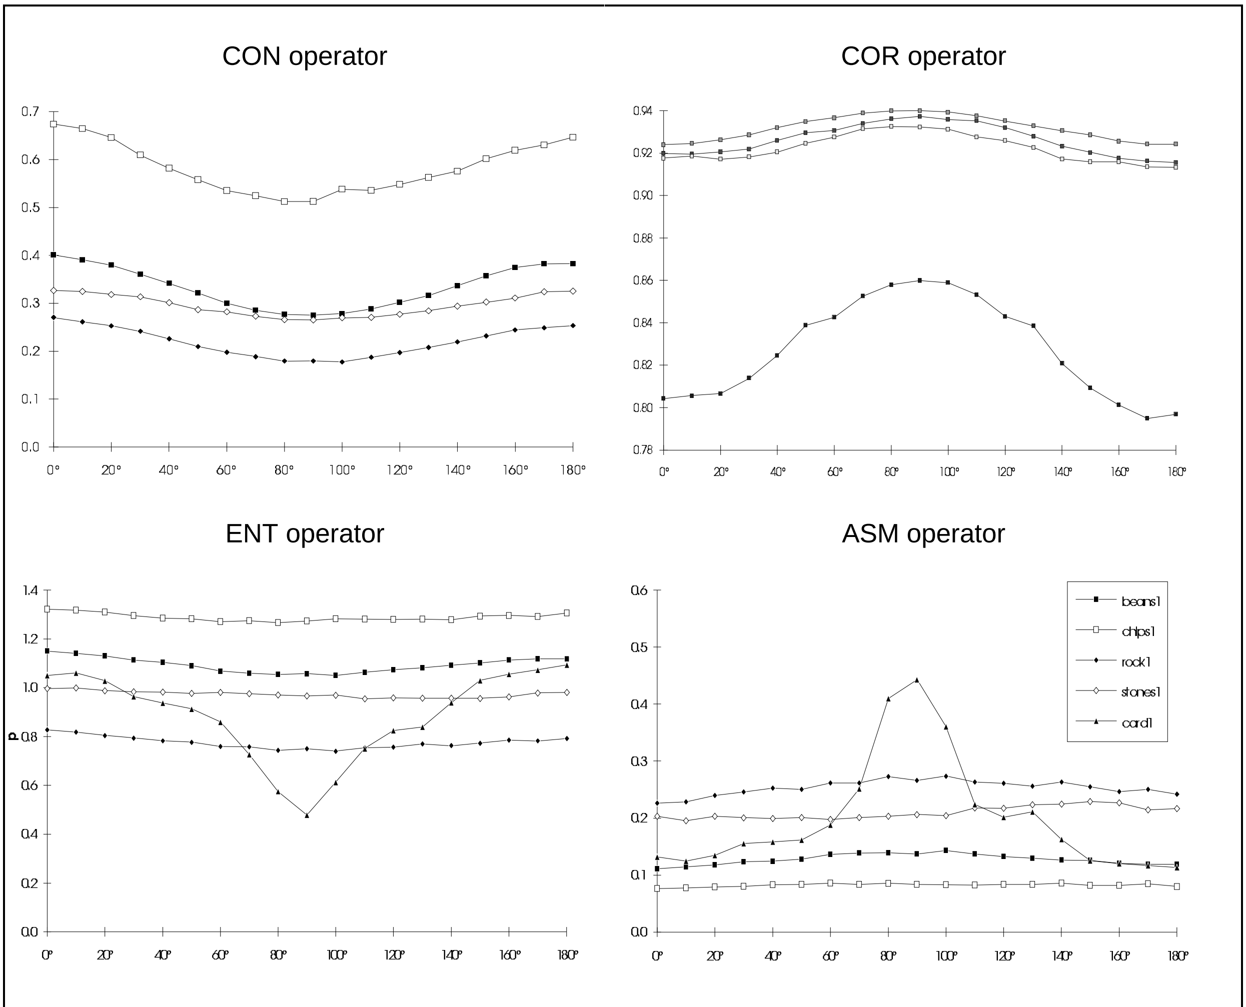

Vertical axes : mean operator output, horizontal axes : illuminant tilt angle

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 θ = 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.

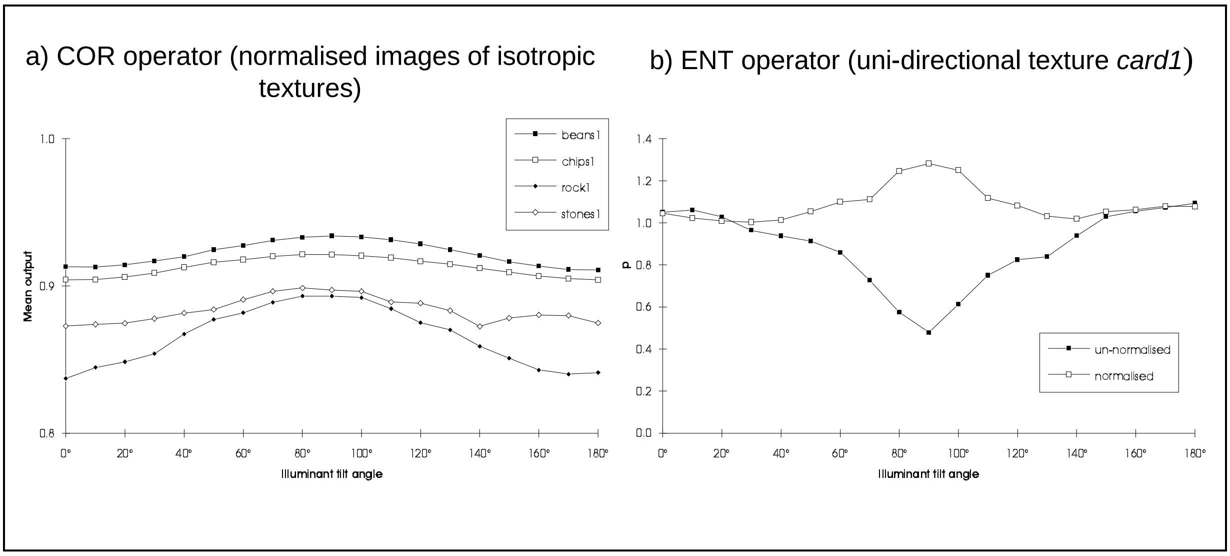

Figure 5.18 - The effect of normalisation on co-occurrence features

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.

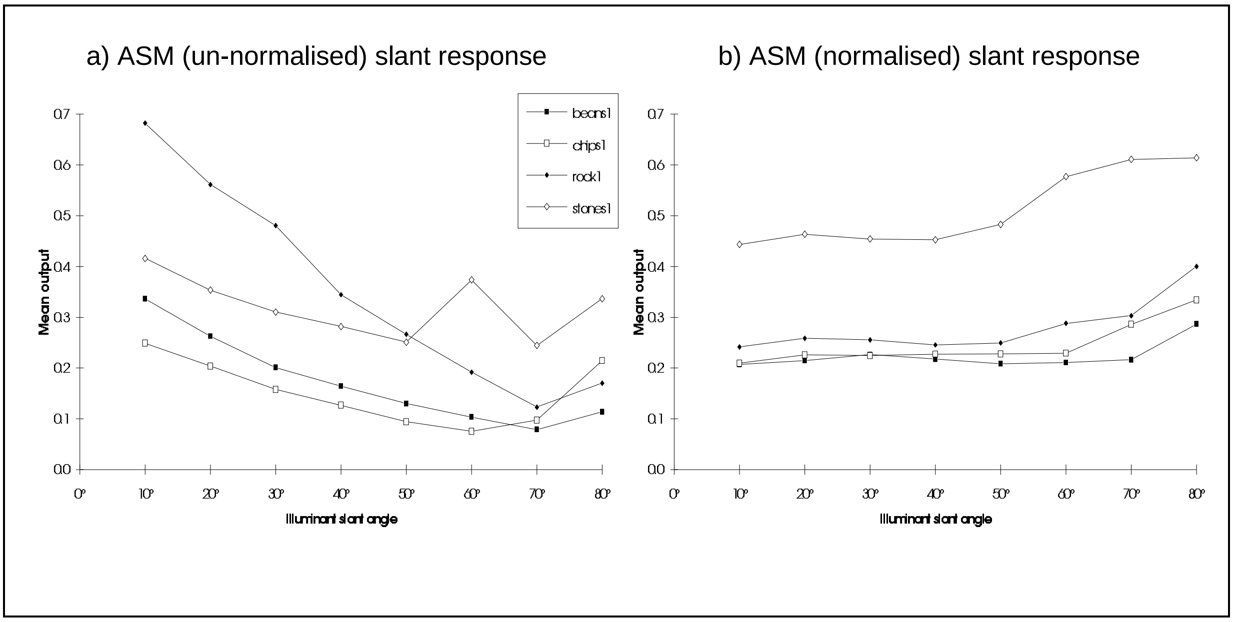

Figure 5.19 - The effect of image normalisation on the ASM *slant response.*

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.

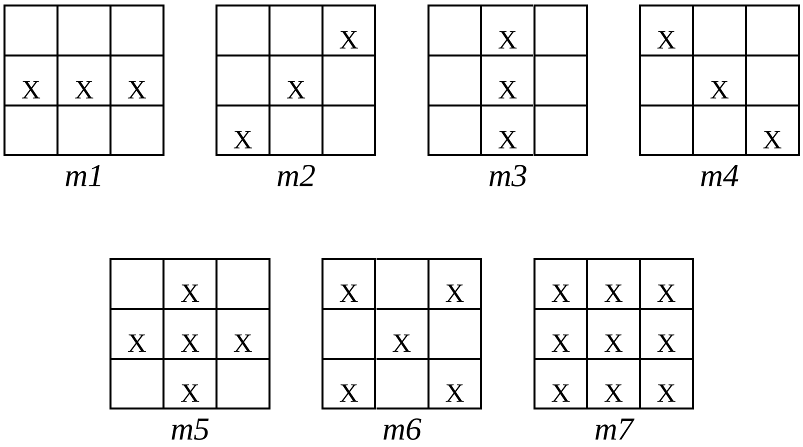

Figure 5.20 - The seven directional masks for Linnett's operator

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.

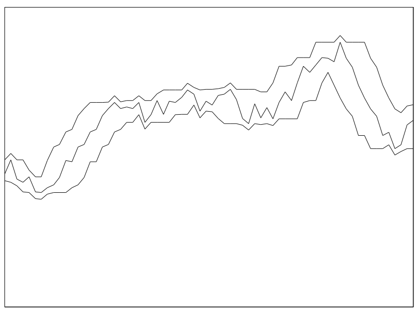

Figure 5.21 - Sections from example surfaces of the blanket method : original surface (middle), upper blanket (top), and lower blanket (bottom)

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 :

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

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

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 :

$$v_n(x,y) = u_n(x,y) - l_n(x,y) \tag{5.37}$$

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 :

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

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

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

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

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

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

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

$$|H_1(\omega)| = 2\sin\omega \tag{5.44}$$

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

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.

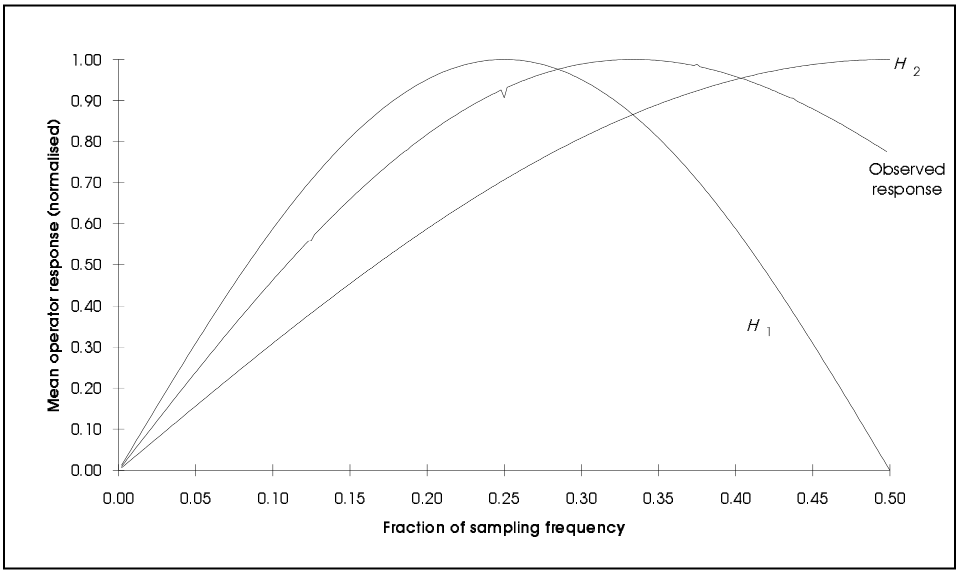

Figure 5.22 - Observed one dimensional frequency response of Linnett's operator, together with the two theoretical cases

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*.

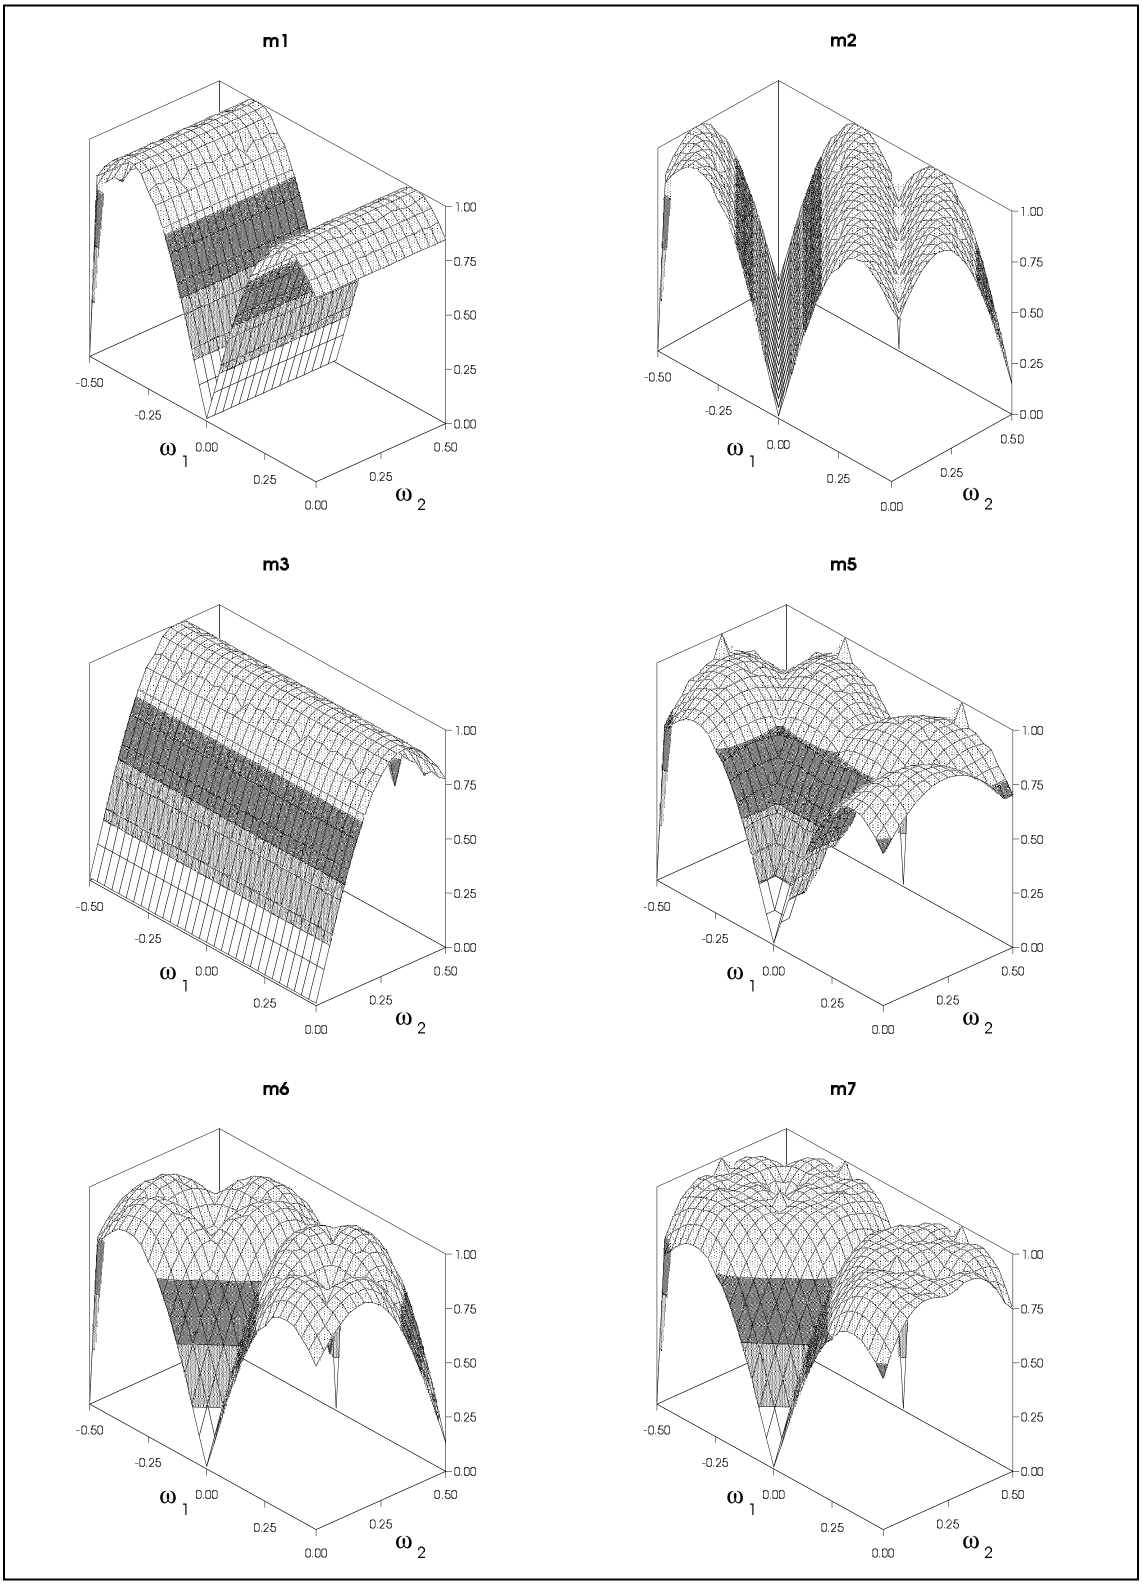

Figure 5.23 - Linnett's operator : observed two dimensional frequency responses

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.

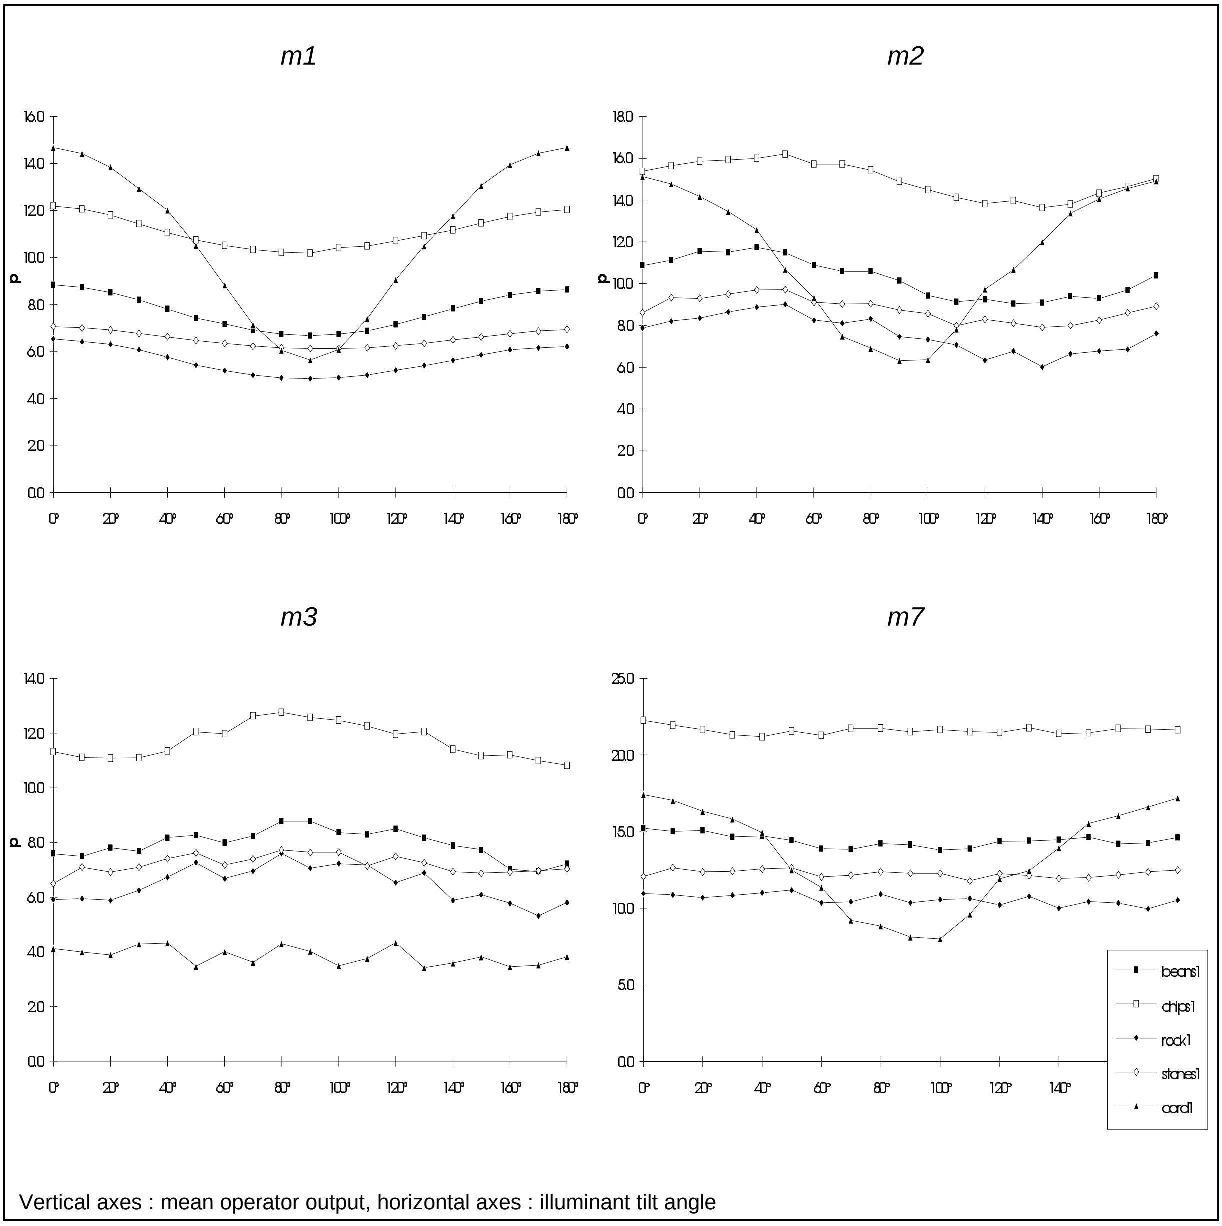

Figure 5.24 - Linnett's operator (un-normalised) : tilt response

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*.

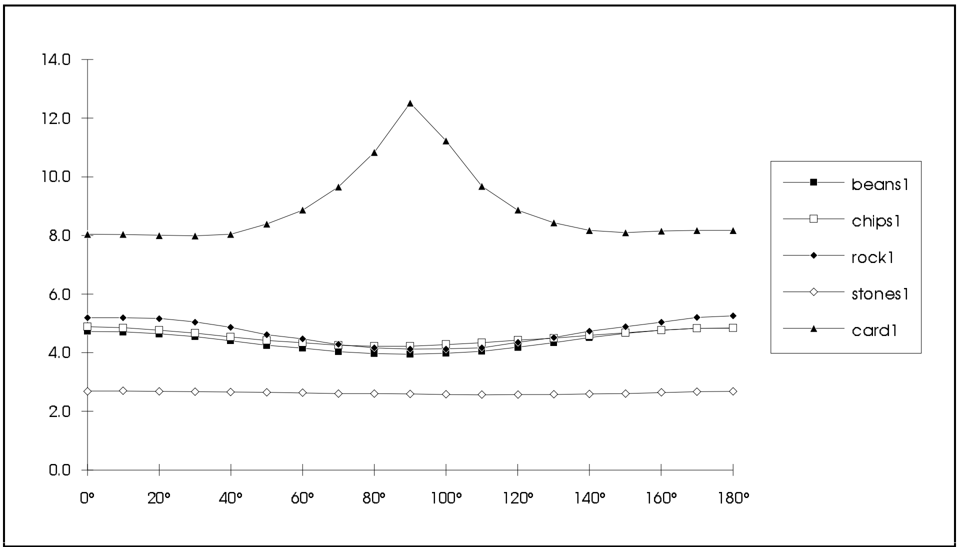

Figure 5.25 - The effect of normalisation on mask m1’s tilt response

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) :

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

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

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

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

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

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

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.

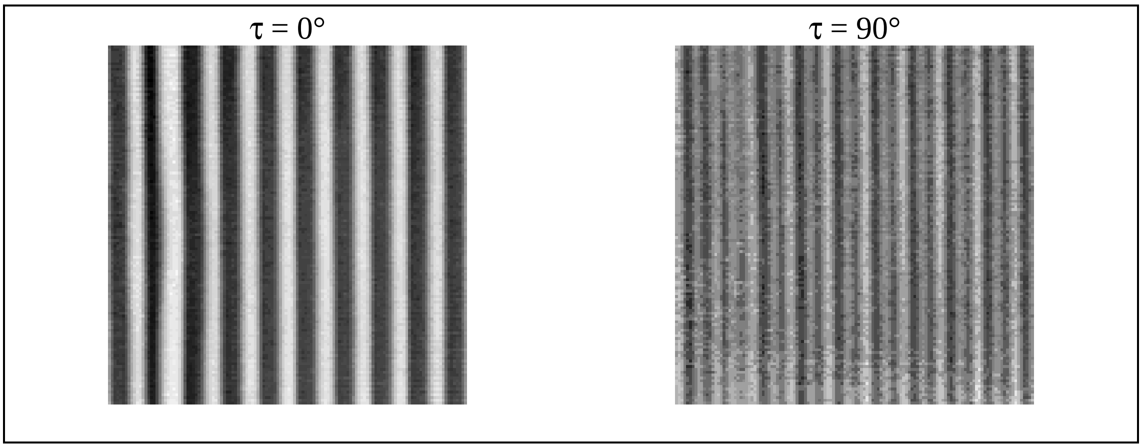

Figure 5.26 - Normalised 128x128 samples of card1 at illuminant tilt angles of 0° and 90° (frequency of corrugations ≅ 0.08 times the sampling frequency)

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.

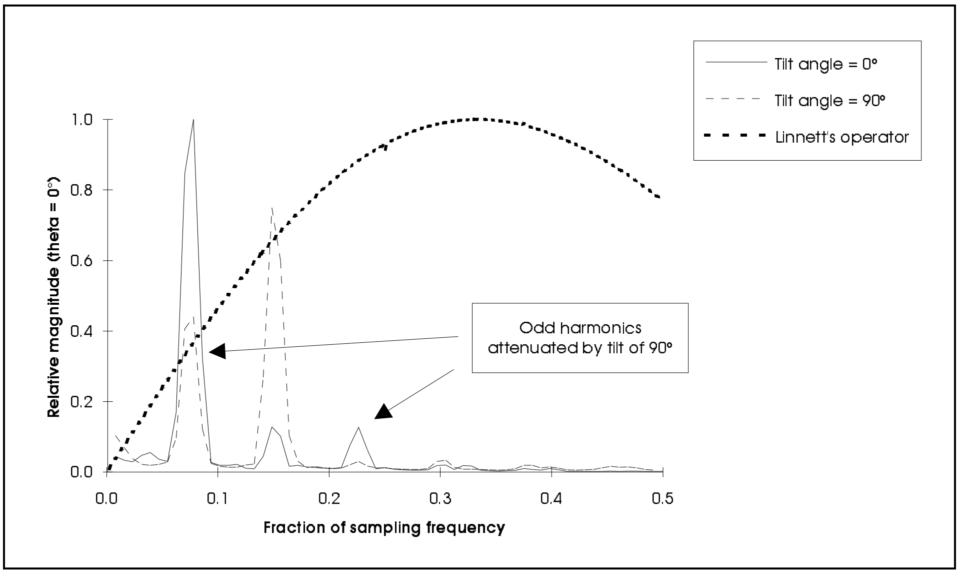

Figure 5.27 - Radial sections of magnitude spectra (at θ = 0°) of card1 and operator m1

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.

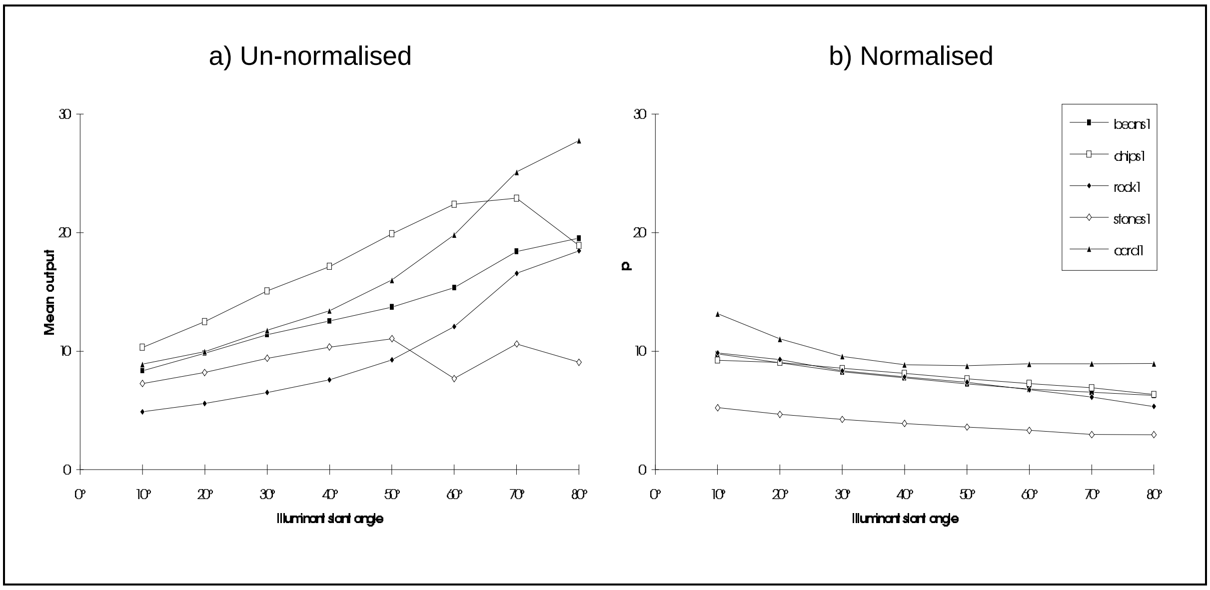

Figure 5.28 - The effect of normalisation on mask m6

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 $\sigma \rightarrow 0$, 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).