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 :

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

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 :

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

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.

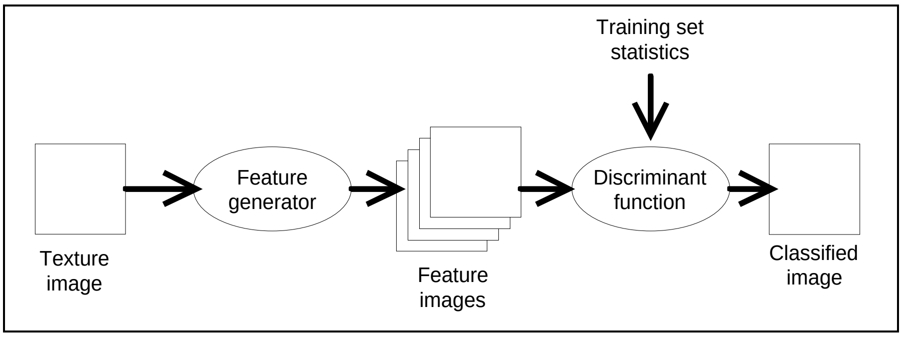

Figure 6.1 - Supervised statistical classification of image texture

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.

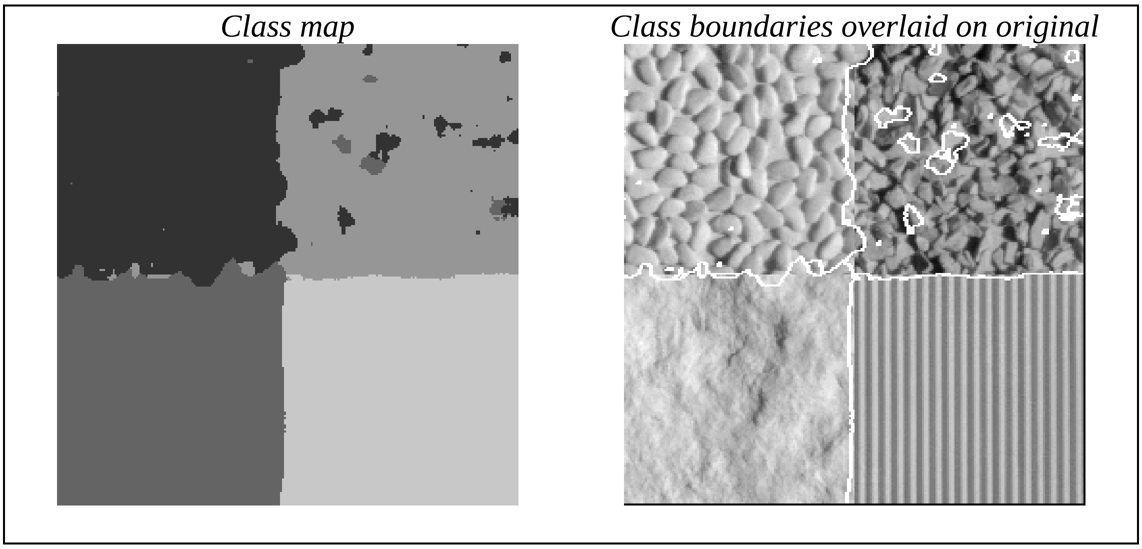

Figure 6.2 - Classification of the four texture image "montage1"

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.

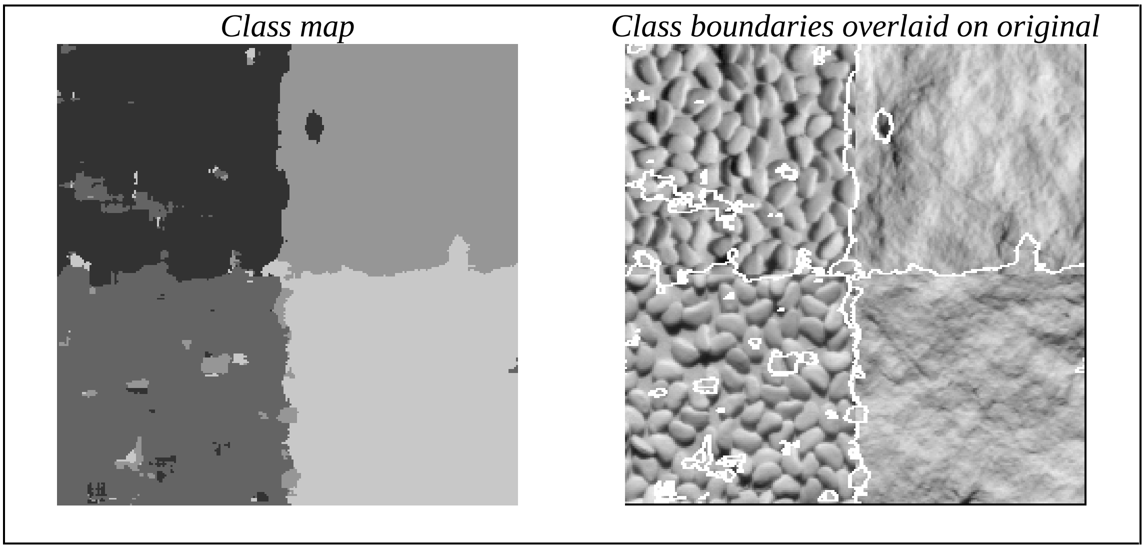

Figure 6.3 - Classification of "montage2" : two physical textures imaged under two illumination conditions

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

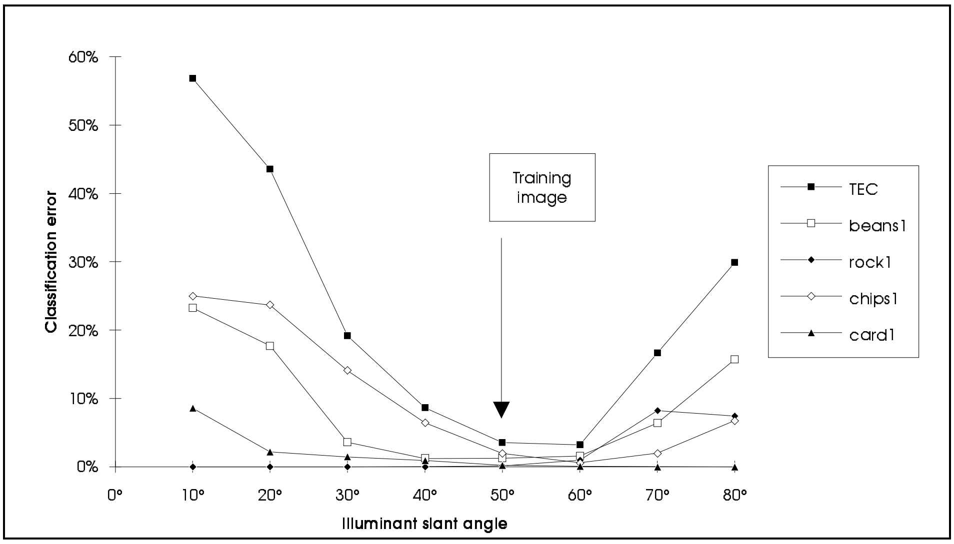

Figure 6.4 - The effect of illuminant slant variation on classifier Laws1

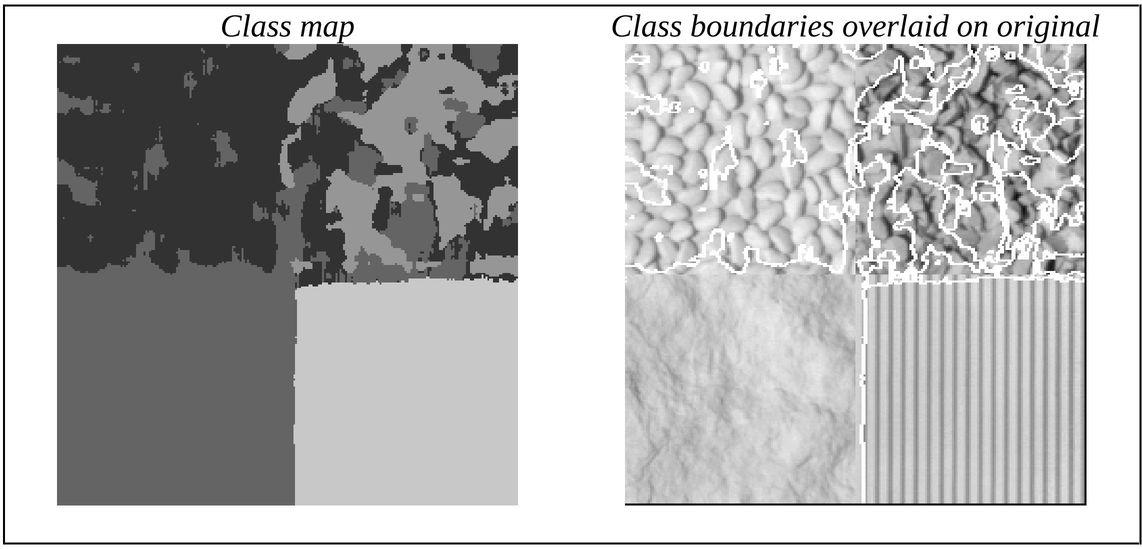

Figure 6.5 - An example of increased failure rate due to variation in illuminant slant (training σ = 50°, test case σ = 30°)

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).

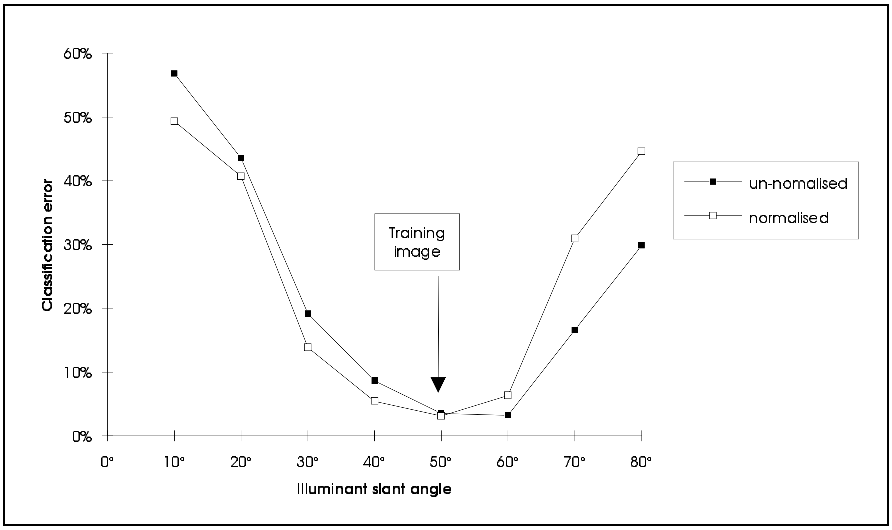

Figure 6.6 - The effect of normalisation on the slant response of the Laws1 classifier (data set : "normalised" montage1)

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 σ = 50° and τ = 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 σ = 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).

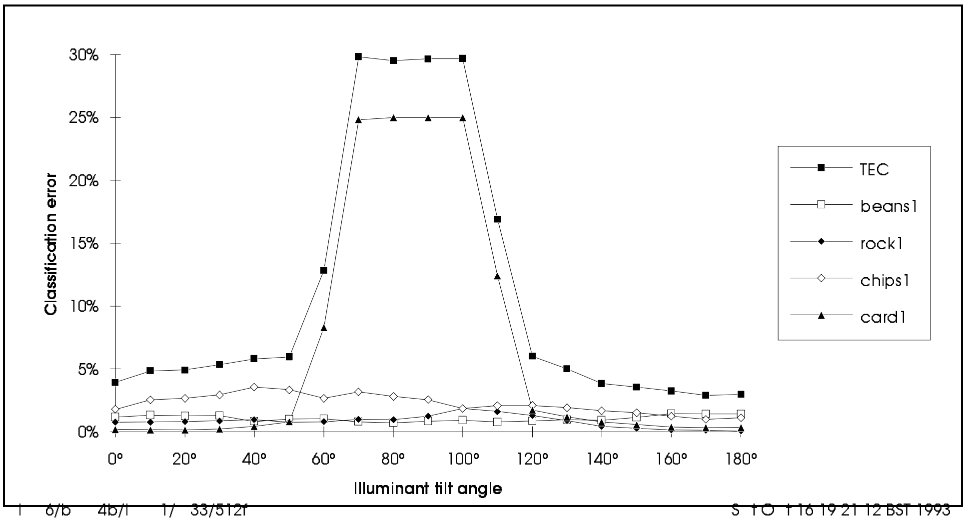

Figure 6.7 - The effect of tilt variation on the classifier Laws1 (data set : montage1)

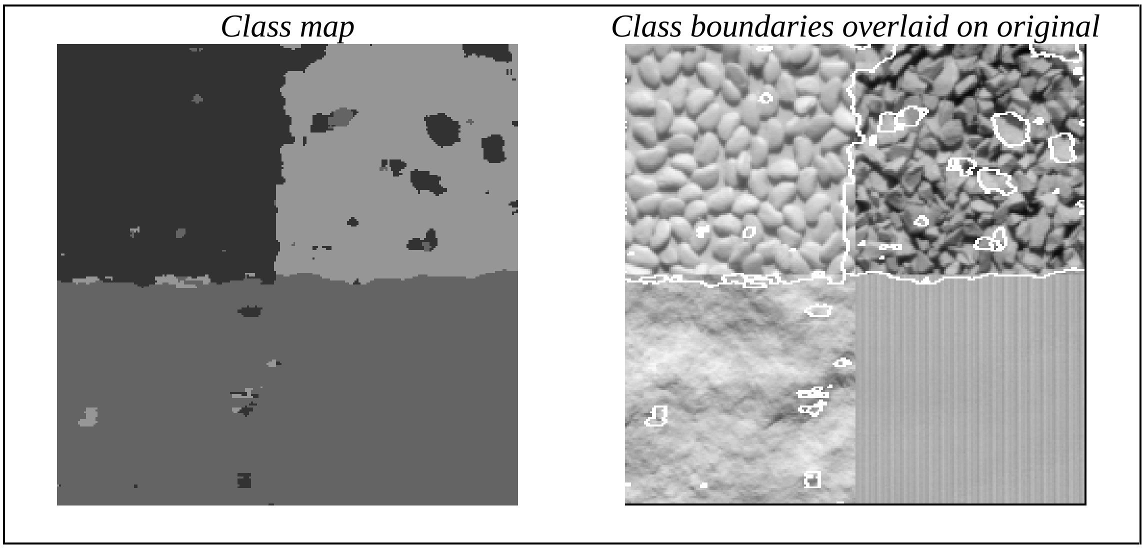

Figure 6.8 - Classification failure at τ = 90° for the Laws1 classifier (data set : montage1)

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.

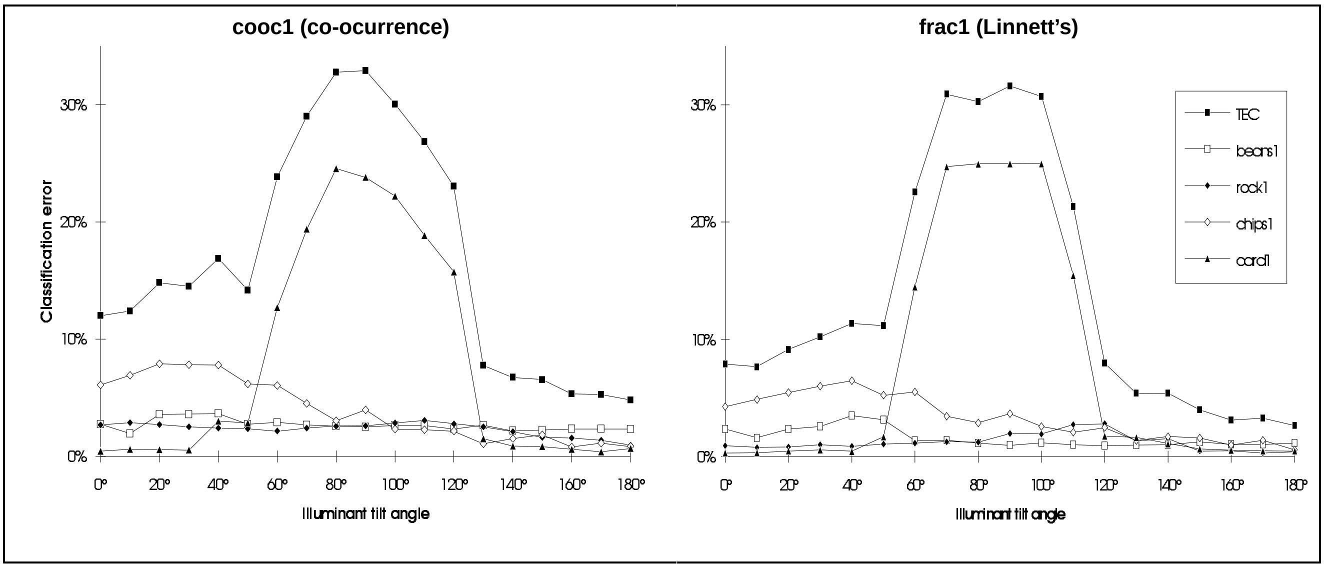

Figure 6.9 - The effect of illuminant tilt variation on "cooc1" and "frac1" classifiers

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).

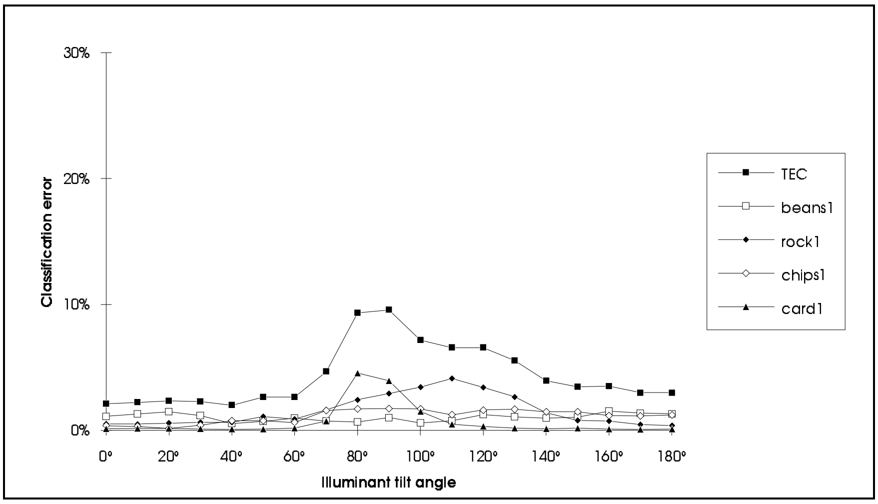

*Figure 6.10 - The effect of normalisation on the previous classification problem (Laws1 classifier; data set : **normalised** montage1)*

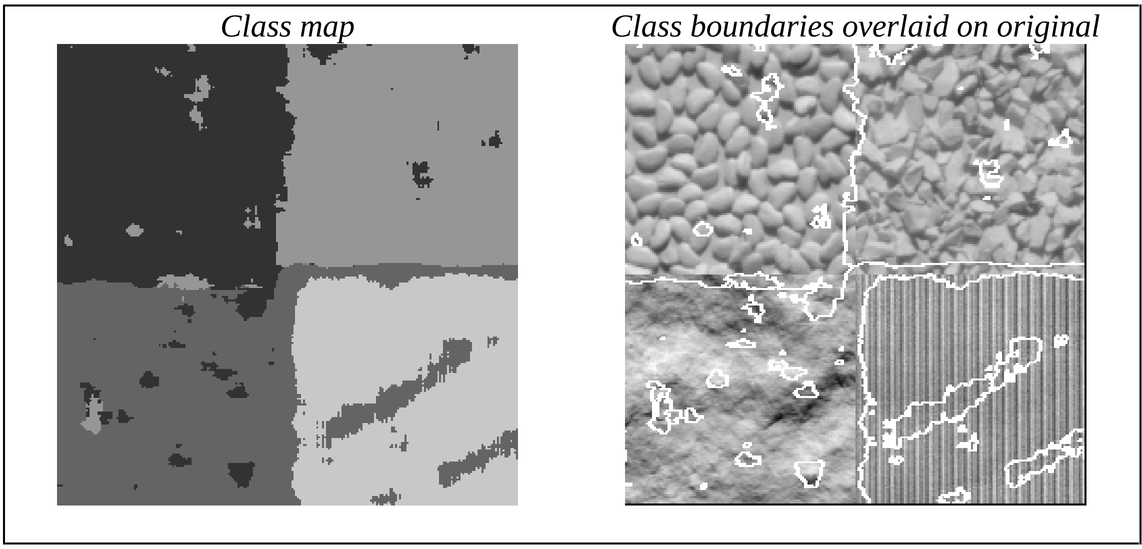

*Figure 6.11 - Classification at τ = 90° (Laws1 classifier, **normalised** montage1)*

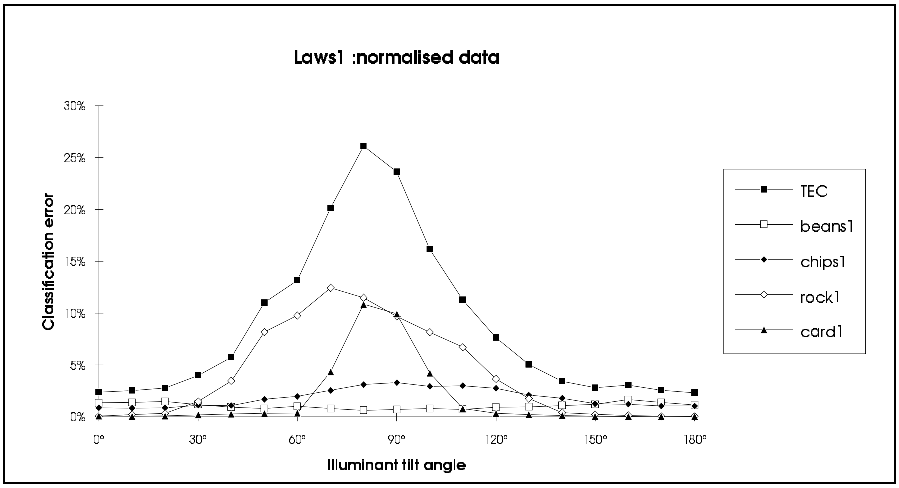

Figure 6.12 - The effect of tilt variation on the Laws1 classifier using normalised images (data set : montage3)

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.

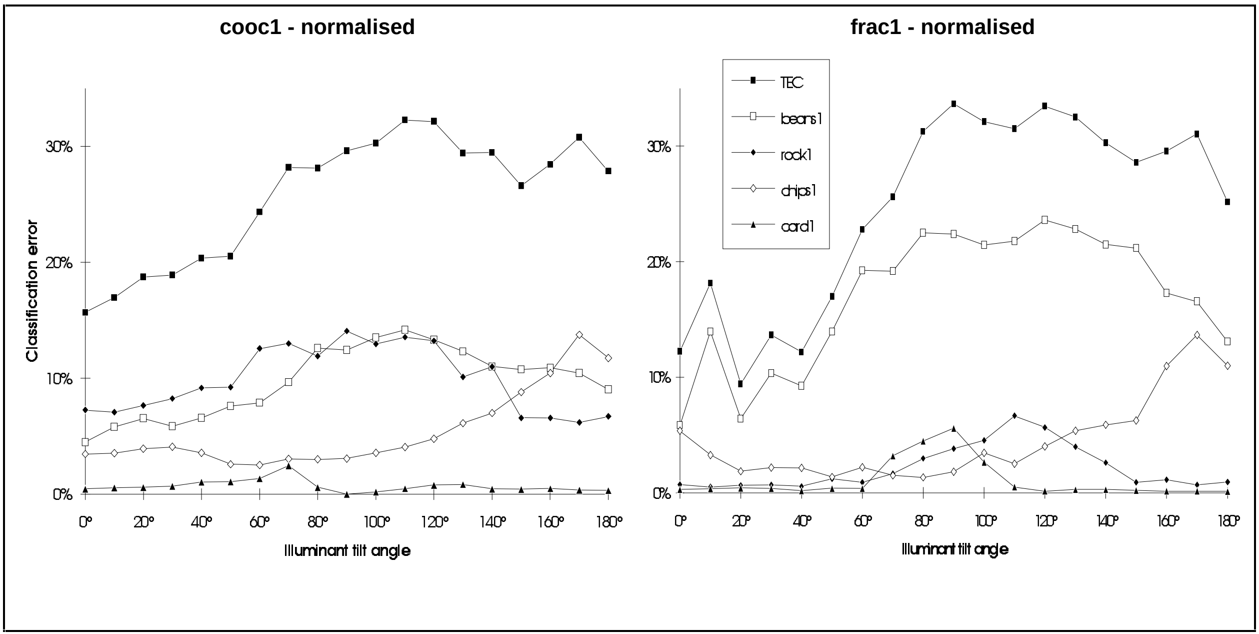

Figure 6.13 - The effects of image normalisation on cooc1 and frac1 classifiers (data set : montage1)

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.

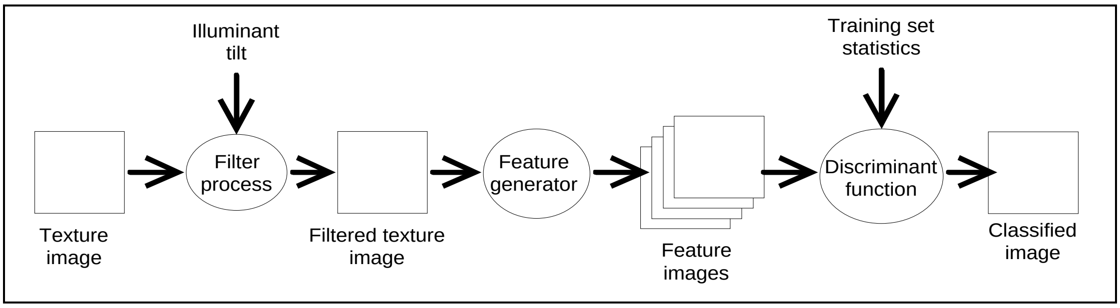

Figure 6.14 - The use of a tilt-compensating filter in the texture classification process

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 τ = 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.

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

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

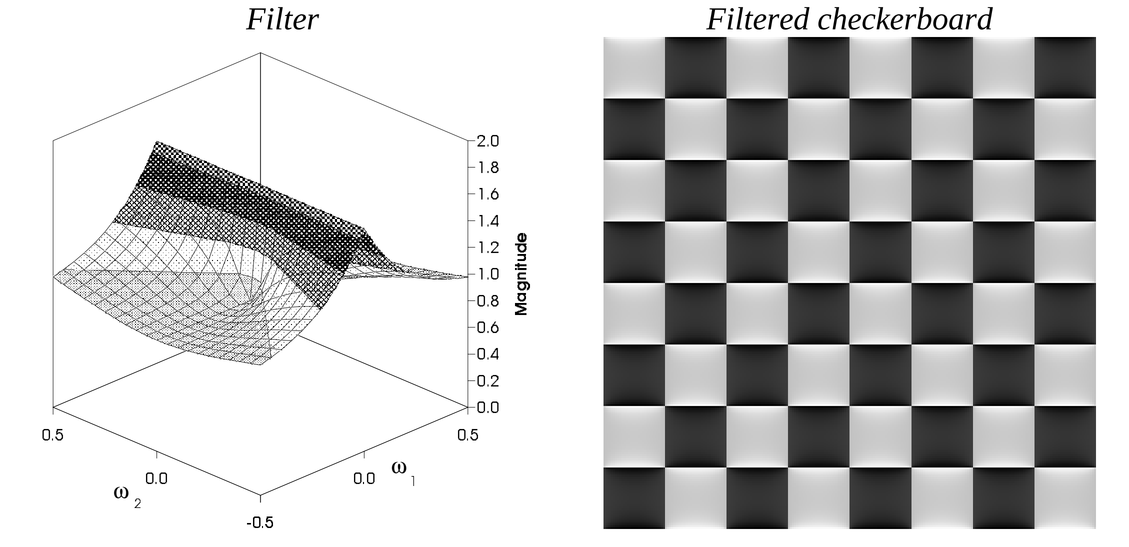

Figure 6.15 - Magnitude frequency response of F1(τ = 0°), and its effect on a checkerboard image.

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".

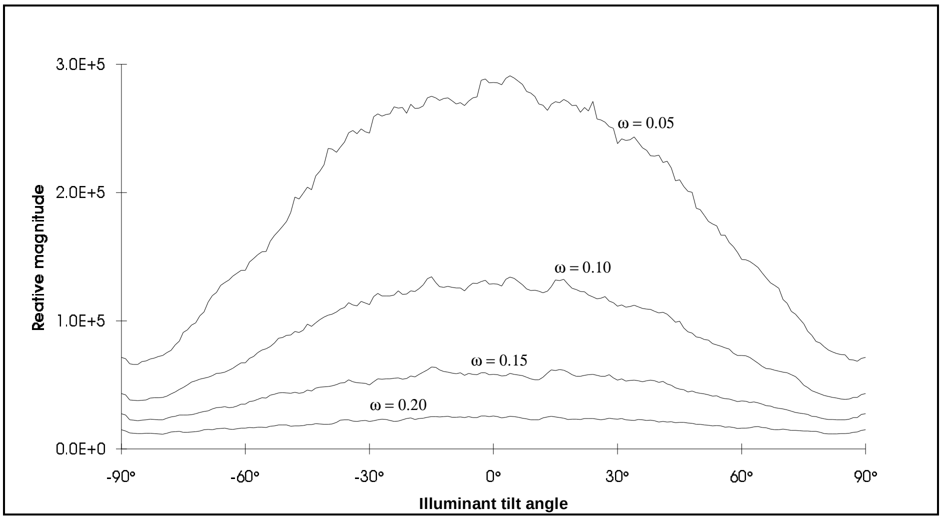

Figure 6.16 - Polar frequency characteristics of rock1 texture (τ = 0°)

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.

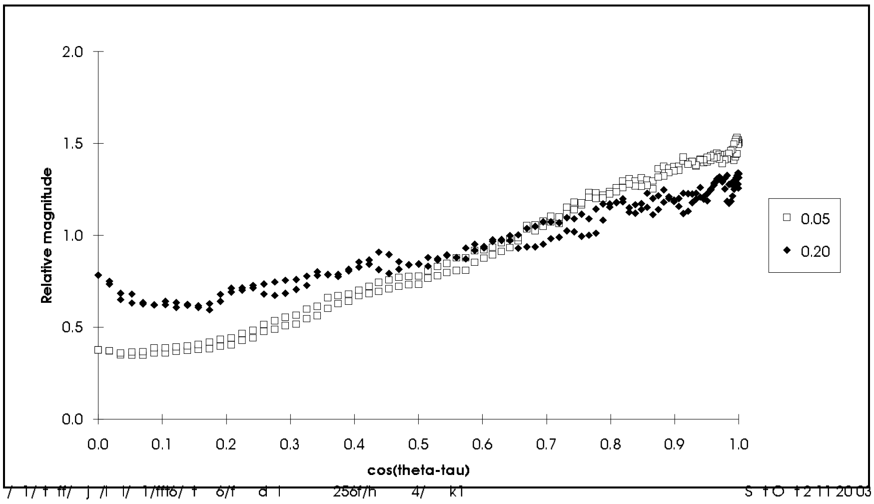

Figure 6.17 - Plot illustrating the $F_I(\omega, \theta) \propto m_\tau \cos(\theta - \tau) + b_\tau$ relationship for $\omega = 0.05$ and 0.20 times the sampling frequency

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

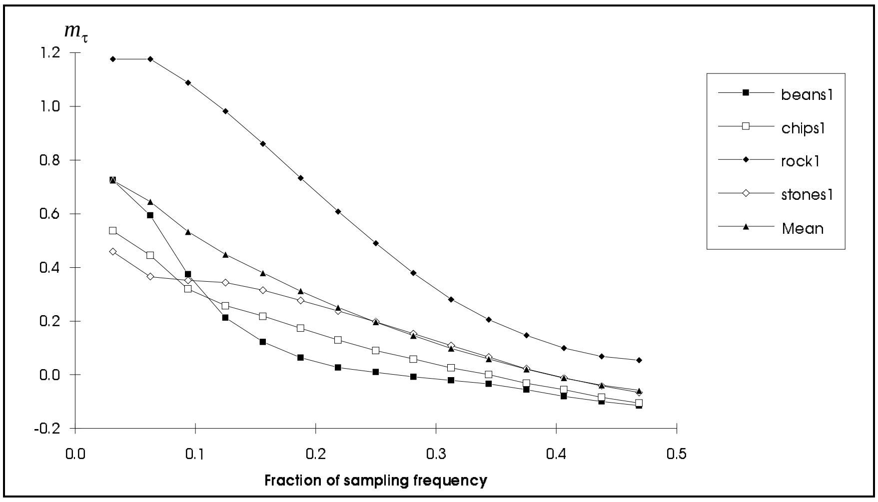

Figure 6.18 - Variation of m_τ with frequency

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.

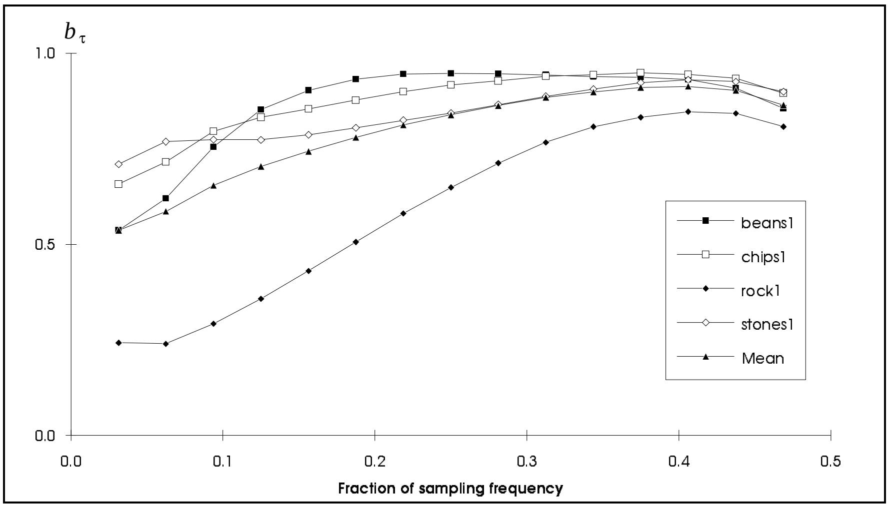

Figure 6.19 - Variation of the parameter b_τ with frequency

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 :

$$m_\tau = -1.8\,\omega/\omega_s + 0.7,\ \ b_\tau = 0.8\,\omega/\omega_s + 0.6 \tag{6.11}$$

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 :

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

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.

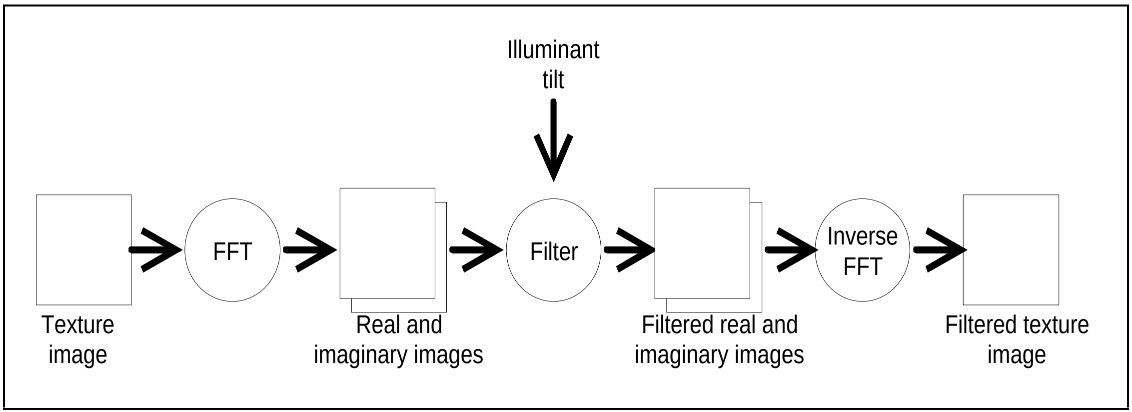

Figure 6.20 - Frequency domain filtering

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*.

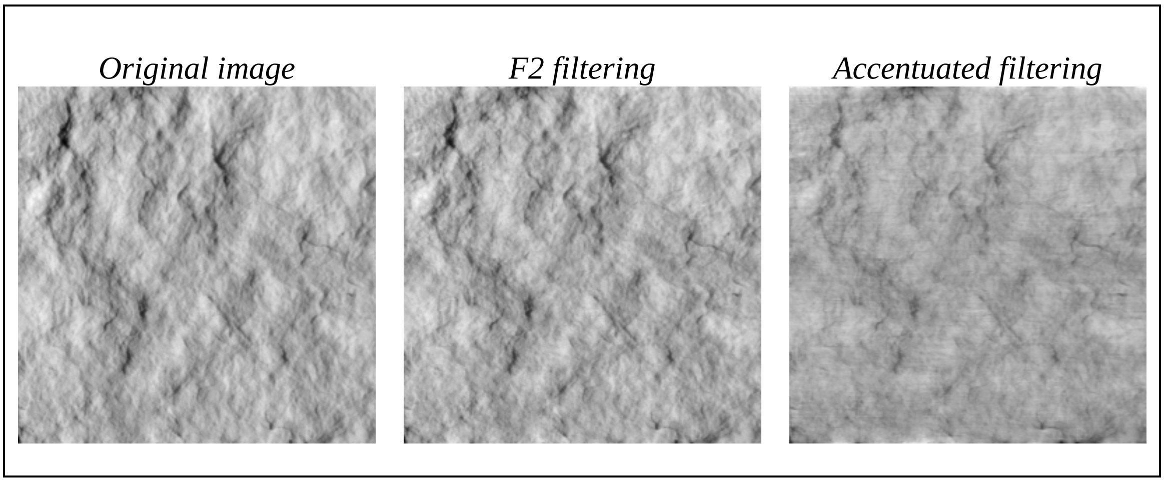

Figure 6.21 - The effect of filter F2(τ = 0°) on the texture "rock1"

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 (τ = 0° and 90°).That is each texture image was processed with the appropriate *F2* compensation filter before application of the L5E5 operator.

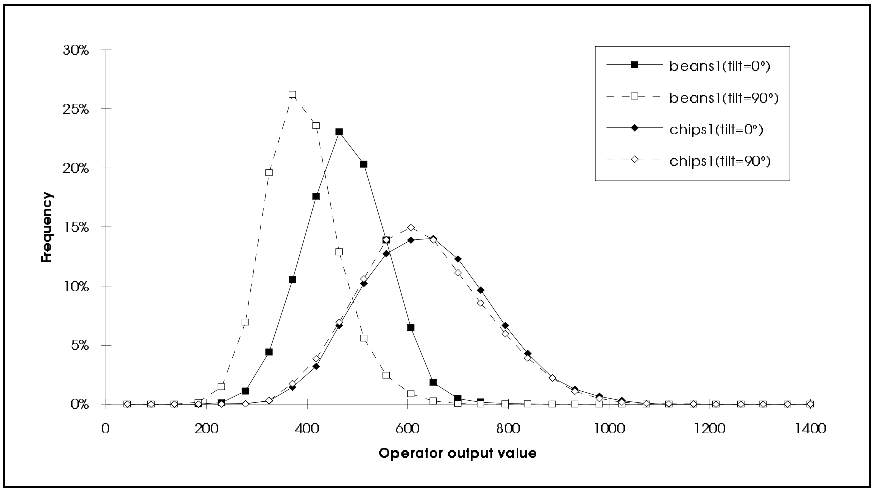

Figure 6.22 - The effect the "F2" filters on tilt behaviour of Laws' L5E5 feature

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.

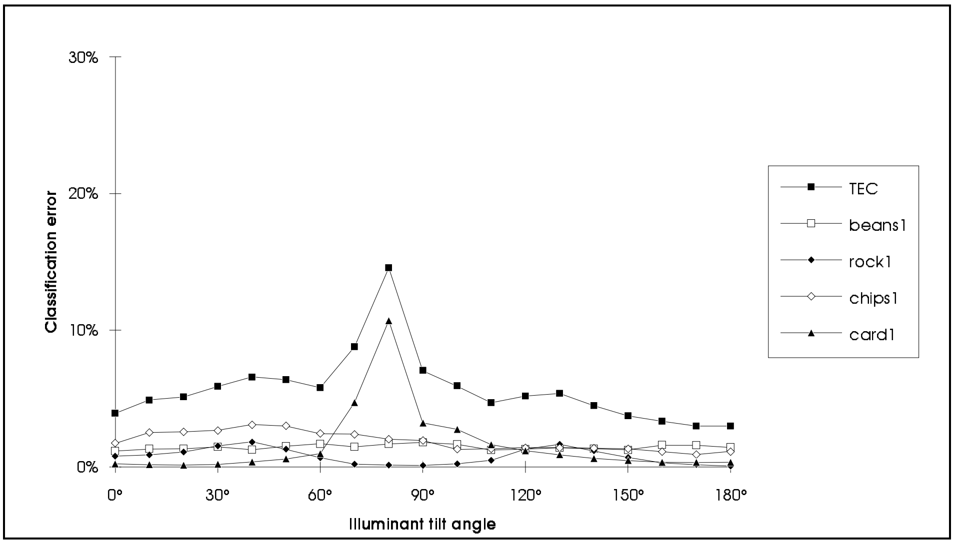

Figure 6.23 - The effect of tilt-compensation on the Laws1 classifier (data set : F2 tilt-compensated montage1)

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°.

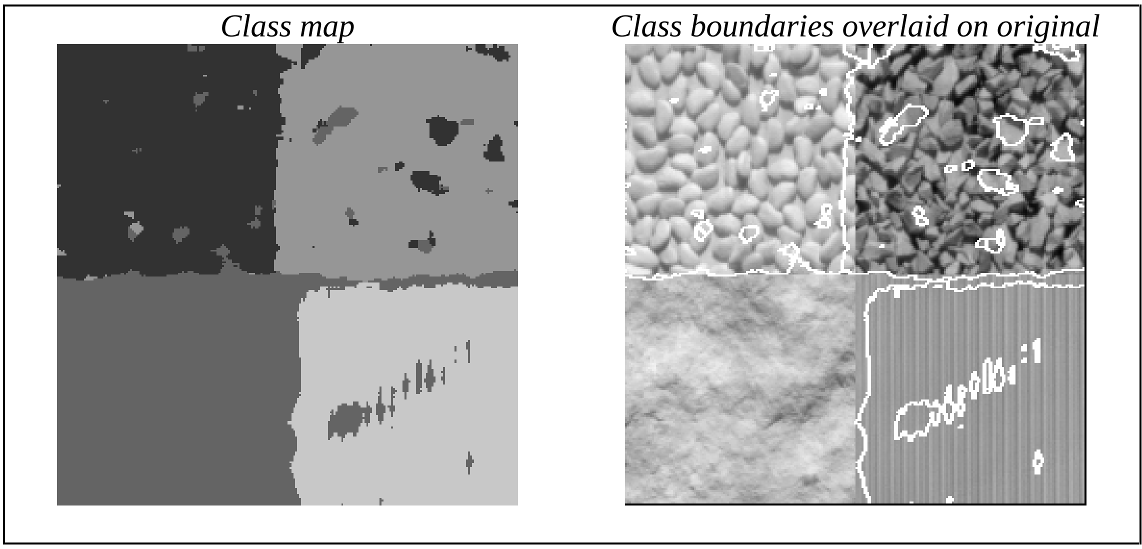

Figure 6.24 - The effect of tilt-compensation on the classification at τ = 90° (Laws1 classifier, data set : F2 tilt-compensated montage1)

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°.

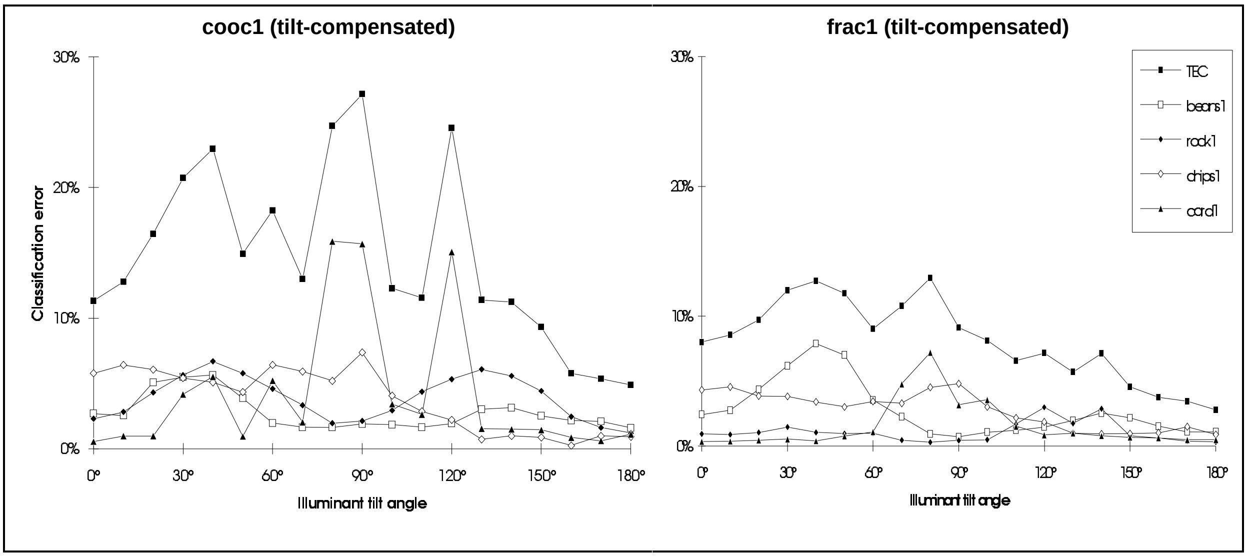

Figure 6.25 - Classification error rates of tilt-compensated cooc1 and frac1 classifiers

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.

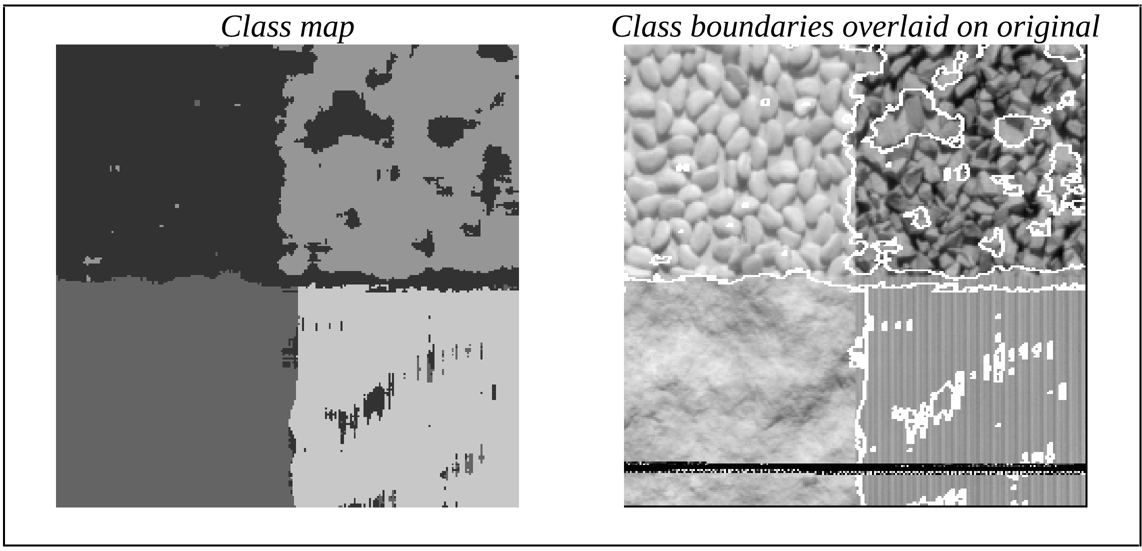

Figure 6.26 - Reduced classification error at τ = 90° (frac1 classifier)

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.

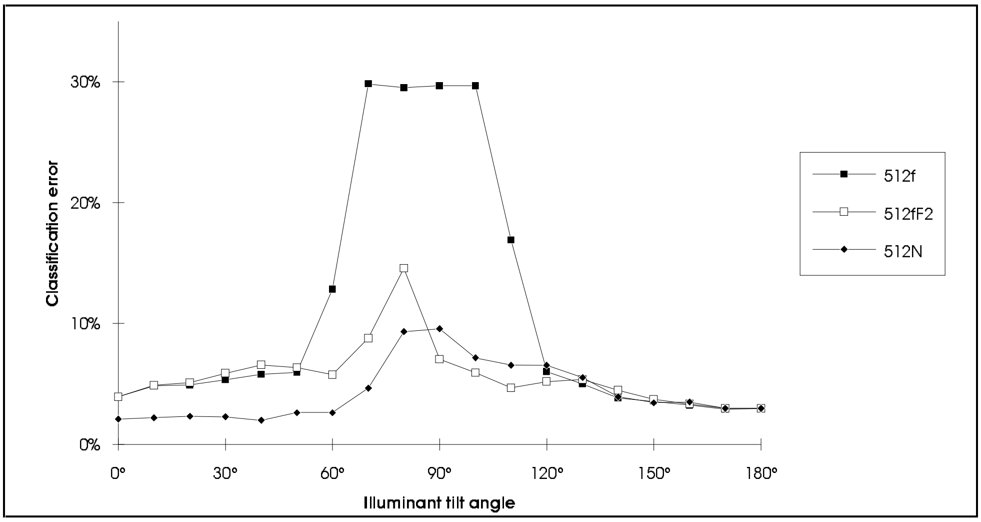

Figure 6.27 - TEC for Laws1 classifier using original(512f), normalised(512N) and tilt-compensated(512fF2) images.

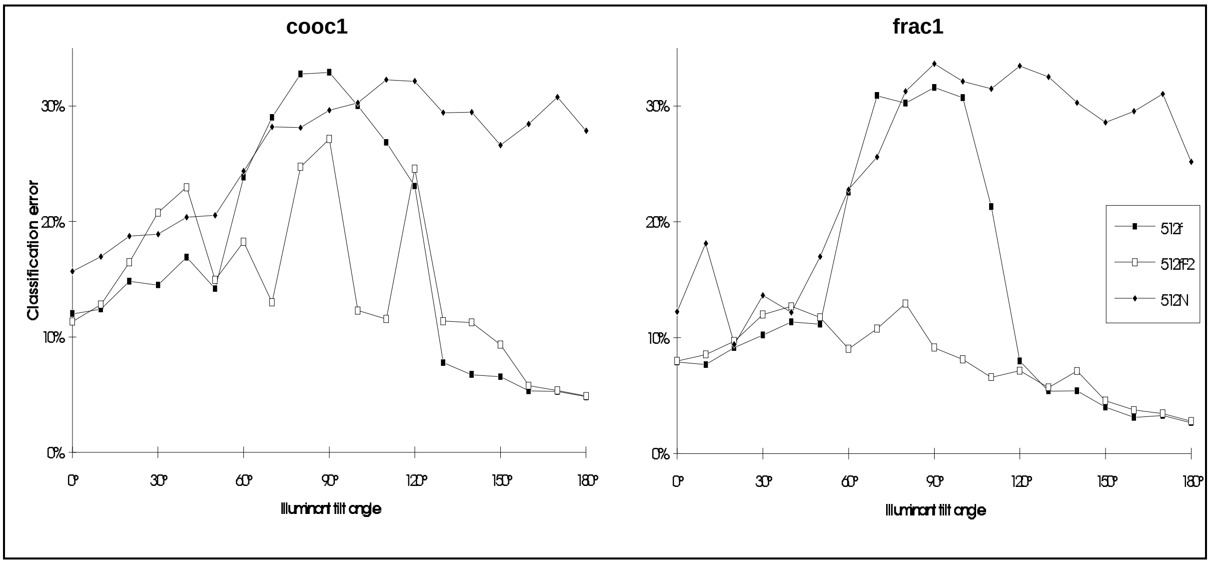

Figure 6.28 - TEC for co-occurrence (left) and Linnett's classifiers (right) using uncompensated, normalised, and tilt-compensated images (512f, 512N and 512f F2 respectively).

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.

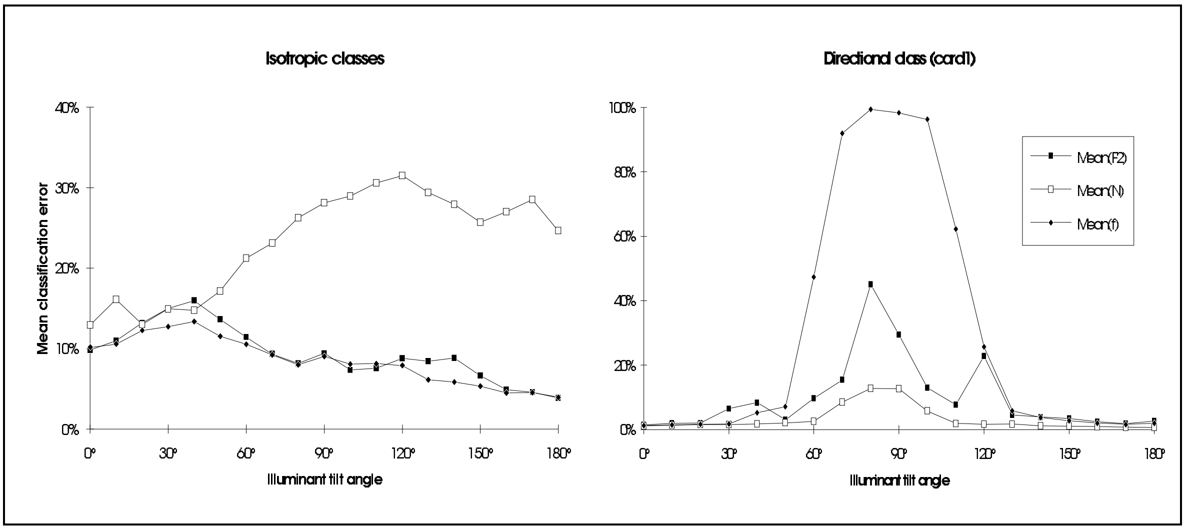

Figure 6.29 - Mean error rates for directional and isotropic texture classes (means calculated from co-occurrence, Laws', and Linnett's features; test set - montage1)

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.

Figures and tables