Chapter 6

Modelling The Classifier Tilt Response

6.1 Introduction

This thesis can be broken into two parts: the first describes the development of a model for rough surface classification, the second part develops techniques to classify surfaces invariant to illuminant tilt. This chapter represents the transition between the analytical and the problem solving phases of this thesis. It has two aims: to model the effect of tilt on the classifier, and to observe the effect of tilt variation on classification accuracy. We also model the classifier developed in Chapter 5; and are therefore concerned with the final stage of the model that describes the process of classification from imaged data set to symbolic representation. We will also measure the effect of tilt variation on classification accuracy. This can therefore be seen as defining the problem that the second part of this thesis is designed to confront.

This chapter consists of four sections. In the first, the tilt response of the texture features will be predicted using theory. The predicted relationships will then be tested using both simulation and experimentation in the second section. The discriminant part of the classifier cannot be modelled in general terms since it depends on the second order interactions between members of the sample set. The next section is designed to extend our analysis and understanding of the tilt effect at the discriminant stage in the classification. The third section of this chapter is designed to give an intuitive understanding of the effect on classification. Using synthetic textures it is possible to obtain an adequate classification using only two features. By using such a small feature set, the discrimination process can be observed on a two dimensional scatter plot—allowing the reader to observe the tilt induced movement of feature clusters across discriminant boundaries with the associated rise in misclassification. Finally, the full classifier is applied to three montages of real textures and the accuracy of the classifier observed as tilt is varied. This chapter will therefore use theory, simulation and experimentation to model the relationship of features with tilt; and use experiments to observe the effect on classification accuracy.

6.2 Modelling The Feature/Tilt Response

In Chapter 5, a classifier was developed. We now set about the formation of an analytical model which encompasses both the classifier and the illumination model of Chapter 3. This will enable us to make predictions as to the effect of illuminant tilt on surface classification.

Chantler made the analogy between the tilt effect and a linear filter [Chantler94b]. We believe this to be a useful model, and combine it with the Gabor filters to form a single linear stage (highlighted in *Figure 6.2.1*), bridging the gap between surface and measure image. We will model the tilt response of the combined filter (sub-section 6.2.1), the power spectrum of the resulting measure images (sub-section 6.2.2), and the first order statistics of the feature images (sub-section 6.2.3).

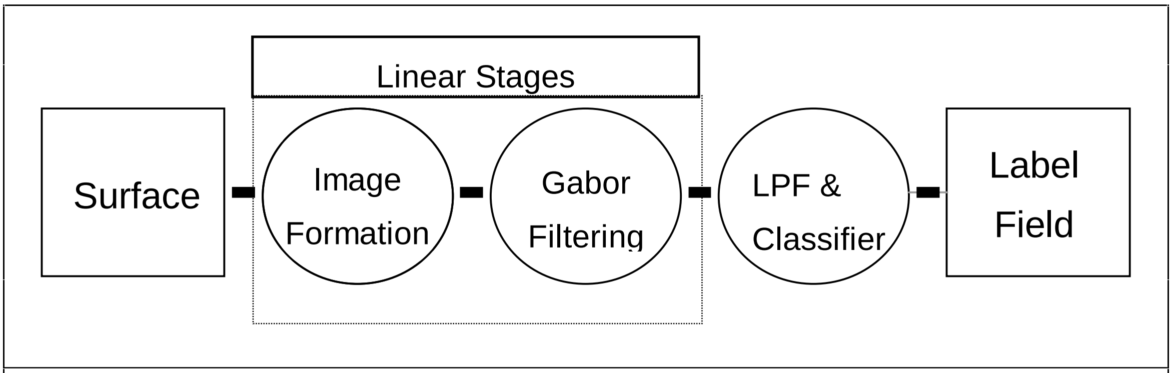

Figure 6.2.1 The Signal to Symbol Chain

6.2.1 Combined Filter Tilt Response

The transfer function relating surface spectra to image spectra may be described as a filter with magnitude response:

$$R(\omega,\theta \mid \tau) = \omega|\cos(\theta-\tau)||k||B(\omega)| \tag{6.2.1a}$$

where $B(\omega)$ is the transfer function of the imaging device, The magnitude response of a Gabor filter oriented in direction ϕ is

$$G\left(u,v\middle|u_o,v_o,\sigma_x,\sigma_y\right)= \quad \exp\left(-0.5\left[\frac{(u-u_o)^2}{\sigma_x{}^2}+\frac{(v-v_0)^2}{\sigma_y{}^2}\right]\right)$$

$$+\exp\left(-0.5\left[\frac{(u+u_o)^2}{\sigma_x{}^2}+\frac{(v+v_0)^2}{\sigma_y{}^2}\right]\right)$$

$$\text{u} = \omega\cos\theta \qquad u_0 = \omega_0\cos\phi$$

$$v = \omega\sin\theta \qquad v_0 = \omega_0\sin\phi$$

In polar form,

$$G(\omega,\theta|\omega_0,\phi,\sigma_u,\sigma_v) = \exp\left(-0.5\left[\frac{(\omega\cos\theta-\omega_0\sin\phi)^2}{\sigma_u{}^2}+\frac{(\omega\sin\theta-\omega_0\cos\phi)^2}{\sigma_v{}^2}\right]\right)$$

$$+\exp\left(-0.5\left[\frac{(\omega\cos\theta+\omega_0\cos\phi)^2}{\sigma_u{}^2}+\frac{(\omega_0\sin\theta+\omega_0\sin\phi)^2}{\sigma_v^2}\right]\right) \quad (6.2.1b)$$

If we combine the rendering response Eq. 6.2.1a with the response of the Gabor filter Eq. 6.2.1b, we obtain the magnitude response of the combined Gabor and rendering filter:

$$Z(\omega,\theta,\tau|\omega_0,\phi,\sigma_u,\sigma_v) = \omega k|\cos(\theta-\tau)|.B(\omega).G(\omega,\theta)+W.G(\omega,\theta) \qquad (6.2.1c)$$

We now numerically integrate equation 6.2.1c over radial frequency, ω, to plot a series of angular response curves in theta for various values of τ, with ω_0=0.125ω_s and ϕ=0°.

$$\int_{\omega=0}^{0.5\omega_s} Z(\omega,\theta,\tau)d\omega \qquad (6.2.1d)$$

where ω_s is the sampling frequency.

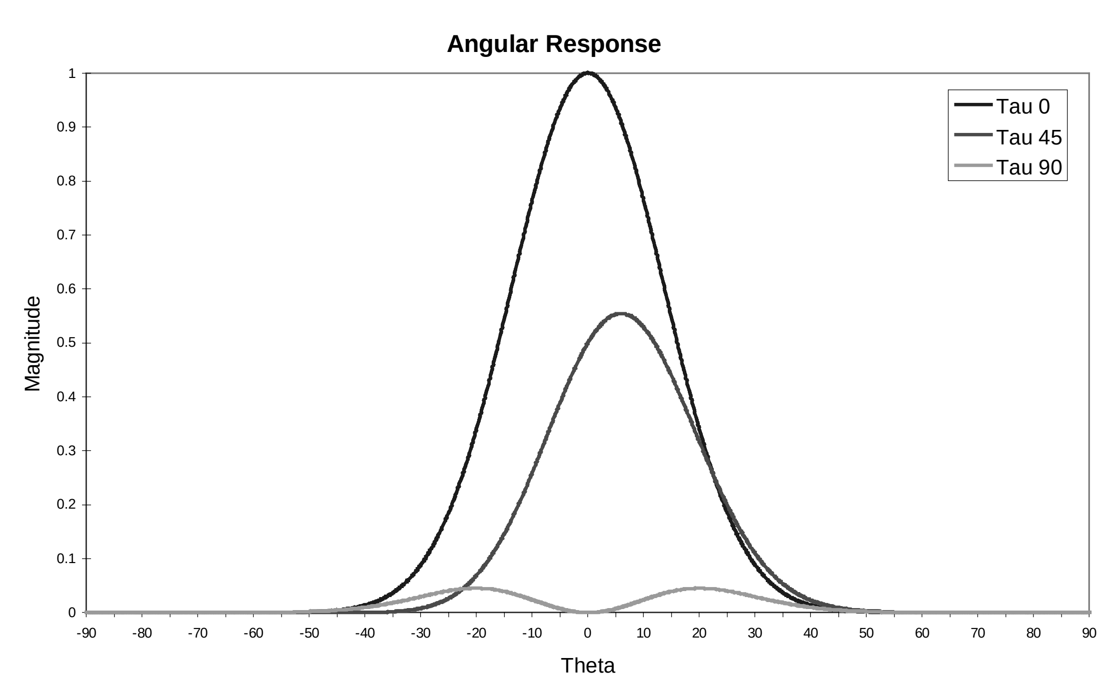

Figure 6.2.2 Angular Transfer Function of Combined Illuminant/Gabor Model

We note three points from this experiment:

1. Firstly, the response of the combined filter is attenuated as the tilt angle τ increases from 0° to 90°. Consequently, the variance of the measure image associated with the filter will fall as the tilt angle is varied through this range.

2. Secondly, the direction of peak sensitivity is shifted in the τ =45° and 90° cases. For an isotropic surface, the directionality of the measure image may no longer be aligned with the directionality of the filter. In the case of a directional surface, the directionality of the measure image will be a function of the surface, illuminant and filter directional properties.

3. Thirdly, even at τ=90°, the filter still gives a significant response due to the bandwidth of the filter orientation. Since we are using a linear model, there is no response at θ=0°, however, there is still a response in other directions due to the two sidelobes present. One consequence of this is that the variance of the measure image for an isotropic image will always be greater than zero. Even with a linear rendering model and in the absence of noise, the feature/tilt response will always be greater than zero.

6.2.2 Modelling The Measure Images

In Chapters 3 and 4, we developed an expression for the power spectral density (PSD) of the data set,

$$I(\omega,\theta \;|\tau) = \omega^2 \cos^2(\theta-\tau)\sin^2\sigma . S(\omega,\theta).|B(\omega)|^2 + W \qquad (6.2.2a)$$

where $S(\omega,\theta)$ is the surface power spectrum.

Each Gabor filter produces a real and an imaginary output image. The PSD of the real and imaginary measure images will be identical, and equal to:

$$D(\omega,\theta|\tau,\omega_0,\phi,\sigma_u,\sigma_v) = I(\omega,\theta).|G(\omega,\theta)|^2$$

Expanding the image term:

$$D(\omega,\theta,\tau|\omega_0,\phi,\sigma_u,\sigma_v) = \omega^2 k^2 \cos^2(\theta-\tau)S(\omega,\theta).|B(\omega)|^2.|G(\omega,\theta)|^2 + N.|G(\omega,\theta)|^2 \quad (6.2.2b)$$

Since the operation is linear, and the original image is assumed to be Gaussian, we may assume the measure images will also have a Gaussian distribution. Integrating (6.2.2b) over both frequency and angle, we will obtain a quantity, which we will call *signal energy*, σ_m^2 , equivalent to the variance of the measure images.

$$\sigma_m^2(\tau) = \int_{\theta=0}^{2\pi}\int_{\omega=0}^{fs} D(\omega,\theta,\tau)d\omega d\theta \qquad (6.2.2c)$$

If we integrate a two dimensional Gabor function, with the parameters used in this thesis, over radial frequency, we obtain a polar plot that is approximately Gaussian with a peak at $\theta=\phi$, *Figure 6.2.3*. We can use this approximation to develop an analytical expression for σ_m^2 by integrating the product of the Gaussian polar response of the Gabor function and the sinusoidal tilt response Eq. 6.2.2d. This is equivalent to the variance of a measure image calculated from the image of an isotropic 1/f surface in the absence of imaging artefacts.

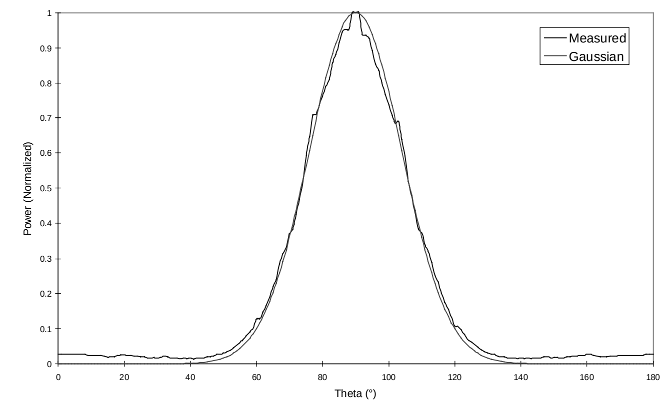

Figure 6.2.3 Comparison of polar response (power) of F25d90 filter with Gaussian function.

$$\sigma_m^2(\theta,\tau,\phi) = \int_{\theta=\phi-\pi}^{\phi+\pi} \cos^2(\theta-\tau)\exp(-0.5\left(\frac{\theta-\phi}{\sigma_p}\right)^2)d\theta \qquad (6.2.2d)$$

For the filters used in this thesis σ_p is less than 20° and the value of the integral is negligable beyond the limits used in 6.2.2d. It is therefore permissable to redefine the range of integration as being from θ=-∞ to ∞.

If we make the substitution $\lambda = \frac{\theta-\phi}{\sqrt{2}\sigma_p}$ then:

$$\theta = \sqrt{2}\sigma_p\lambda + \phi \qquad \frac{d\lambda}{d\theta} = \frac{1}{\sqrt{2}\sigma_p}$$

when θ=-∞ then λ=-∞ and when θ=-∞ then λ=∞

Therefore when integrating between θ=-∞ to ∞, the equivalent limits for integration are λ=-∞ to ∞.

$$\sigma_m^2(\theta,\tau,\phi) = \int_{\lambda=-\infty}^{\infty} \sqrt{2}\sigma_p \cos^2\left(\sqrt{2}\sigma_p\lambda + \phi - \tau\right)\exp\left(-\lambda^2\right)d\lambda$$

$$\sigma_m^2(\theta,\tau,\phi) = \frac{\sigma_p}{\sqrt{2}} \int_{\lambda=-\infty}^{\infty} \left(1+\cos\left(\sqrt{8}\lambda\sigma_p + 2(\phi-\tau)\right)\right)\exp\left(-\lambda^2\right)d\lambda$$

Let $\alpha = \sqrt{8} \qquad \beta = 2(\phi - \tau)$

$$\sigma_m^2(\theta,\tau,\phi)=\frac{\sigma_p}{\sqrt{2}}\int_{\lambda=-\infty}^{\infty}\left(1+\cos(\alpha\lambda+\beta)\right)\exp\left(-\lambda^2\right)d\lambda$$

Using the general results:
$\int_{\lambda=-\infty}^{\infty}\exp\left(-\lambda^2\right)d\lambda=\sqrt{\pi}$
$\int_{\lambda=-\infty}^{\infty}\cos(\alpha\lambda-\beta)\exp\left(-\lambda^2\right)=\sqrt{\pi}\cos(\beta)\exp\left(\frac{-\alpha}{4}\right)$

$$\sigma_m^2=\frac{\sigma_p\sqrt{\pi}}{\sqrt{2}}\left(1+\cos(\beta)\exp\left(-\frac{\alpha}{4}\right)\right)$$

$$\sigma_m^2(\tau,\phi)=\frac{\sigma_p\sqrt{\pi}}{\sqrt{2}}\left(1+\left(2\cos^2(\phi-\tau)-1\right)\exp\left(-\frac{\alpha}{4}\right)\right)$$

If we gather the angular variables, we obtain an expression for the variance of the measure image, which is a function of the filter orientation and the illuminant tilt:

$$\sigma_m^2(\phi,\tau)=a\cos^2(\phi-\tau)+b \qquad (6.2.2e)$$

where *a* and *b* are constants.

Consequently, we predict that the variance of the measure image will vary with the square of the cosine of the angle between the illuminant vector and the direction of maximum sensitivity of the Gabor filter. We can clearly observe this result if we integrate equation 6.2.2c numerically. The resulting function can then be sampled at values of ϕ corresponding to our filter set and plotted as a fuction of tilt, with ω_0=0.125ω_s.

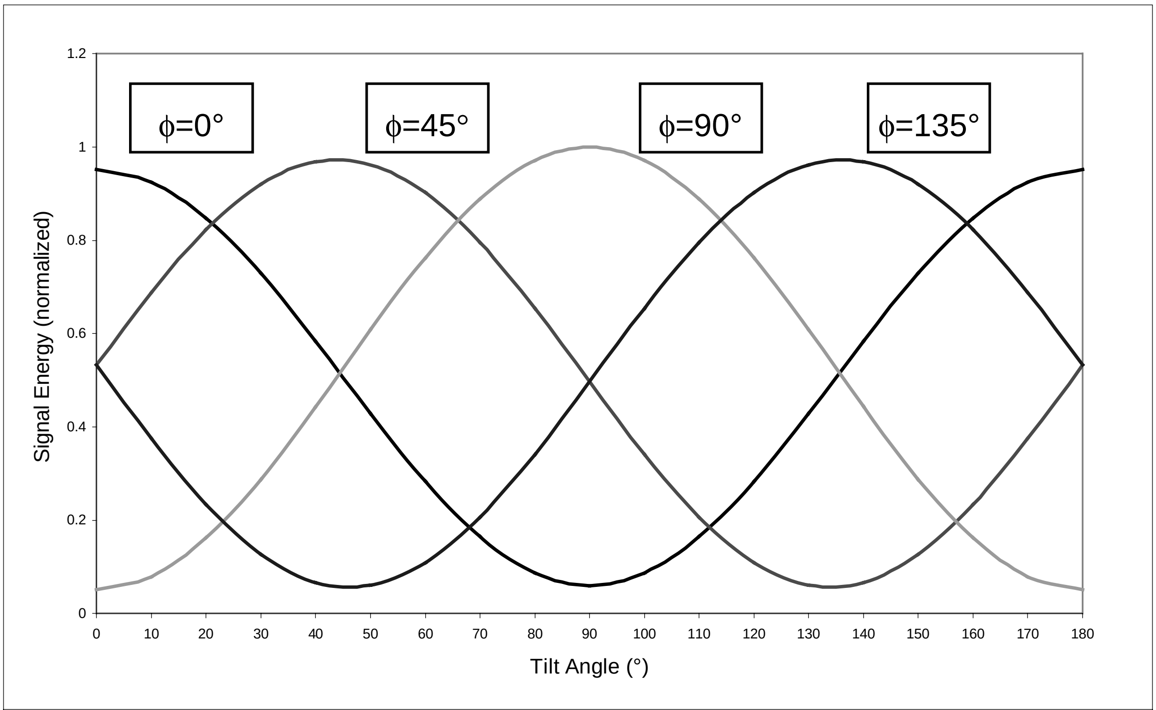

Figure 6.2.4 Numerical estimate of measure variance for an isotropic fractal surface.

The result of the numerical integration and evaluation are plotted against $cos^2(\phi-\tau)$ in *Figure 6.2.5*. In all cases a near linear relationship exists. The function $a\cos^2(\phi-\tau)+b$ is also plotted, where the parameters *a* and *b* are the fitted using least squares to give the best fit to the integration results. In the case of the ϕ=0° and ϕ=90° curves, the best fit is visually indistinguishable from the evaluated curves and is not plotted. In the case of the ϕ=45° and ϕ=135° the numerically evaluated curves exhibit a small amount of hysteresis—presumably due to numerical approximations. The hysteresis is symmetrical about the best fit curve. The numerical results agree well with the analytical predictions, (6.2.2e).

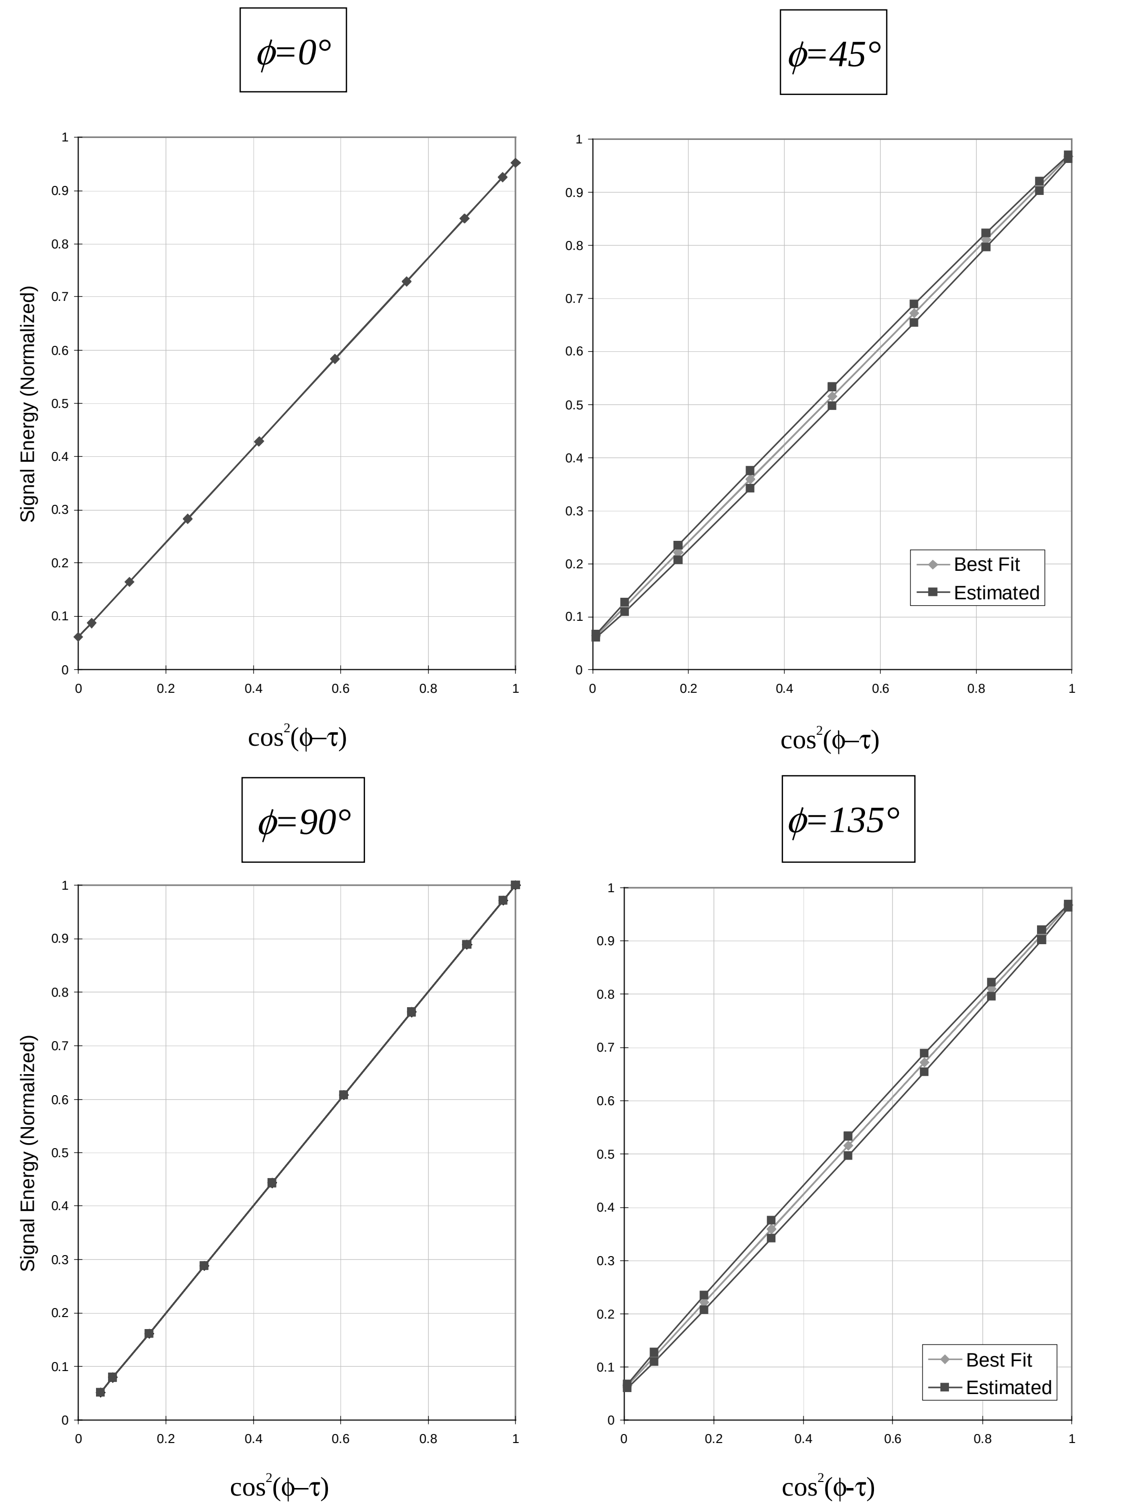

Figure 6.2.5 The numerical integration of Eq. 6.2.2c for various values of ϕ.

6.2.3 Feature Image Statistics

Assuming both the measure images to be zero mean Gaussian with standard deviation σ_m, the magnitude image will have a Rayleigh distribution, Eq. 6.2.3a, parameterised by σ_m:

Rayleigh distribution: $$p(f)=\frac{f}{\sigma_m^{\;2}}\exp\left(\frac{-f^2}{2\sigma_m^{\;2}}\right) \tag{6.2.3a}$$

with mean, $$\mu_f=\sqrt{\frac{\pi}{2}}\sigma_m \tag{6.2.3b}$$

and standard deviation, $$\sigma_f=\sqrt{2-\frac{\pi}{2}}\sigma_m \tag{6.2.3c}$$

As was noted in the previous chapter, it is necessary to low pass filter the magnitude image. Due to the earlier non-linear operation of calculating the magnitude image from the quadrature images, it is not possible to model this in the frequency domain in the context of the earlier results. Nevertheless, the mean response of the magnitude will be largely unaffected by the low pass filtering, merely scaled by the sum of the filter weights. Since the mean of the distribution is a linear function of σ_m, (Eq. 6.2.3b), we predict from Eq. 6.2.2e that the feature mean will vary with a relationship of the form:

$$\mu_f(\tau)=a\left|\cos(\phi-\tau)\right|+b$$

Although we cannot predict the value of the standard deviation, we can hypothesise its relationship with tilt angle. The low pass filter is isotropic, and our (linear) model of rendering induced directionality is independent of radial frequency (if we ignore imaging artefacts). Consequently, it would be reasonable to suggest that the standard deviation of the filter shares the tilt characteristics of the magnitude image, since the standard deviation of the magnitude image is also a linear function of σ_m (Eq. 6.2.1k). We therefore hypothesise that both the mean and standard deviation of the feature follow a relationship of the form:

$$\sigma_f=a\left|\cos(\phi-\tau)\right|+b$$

We shall now compare this prediction with the filter output obtained from both simulation and experiment.

6.3 Testing the theoretical predictions

6.3.1 Verification by Simulation

The figure below shows the measured mean and standard deviation of a feature derived from a single ϕ=0° filter applied to an synthetic, isotropic, fractal surface. Both the mean and the standard deviation are scaled to give a maximum of unity. We also plot two additional waveforms for comparison, the absolute value of the cosine and the square of the cosine. The parameters *a* and *b* have, in both cases, been estimated using a least squares fit to the average of the mean and the standard deviation values.

$$O(\tau) = a|\cos(\tau)| + b$$

$$O(\tau) = a\cos^2(\tau) + b$$

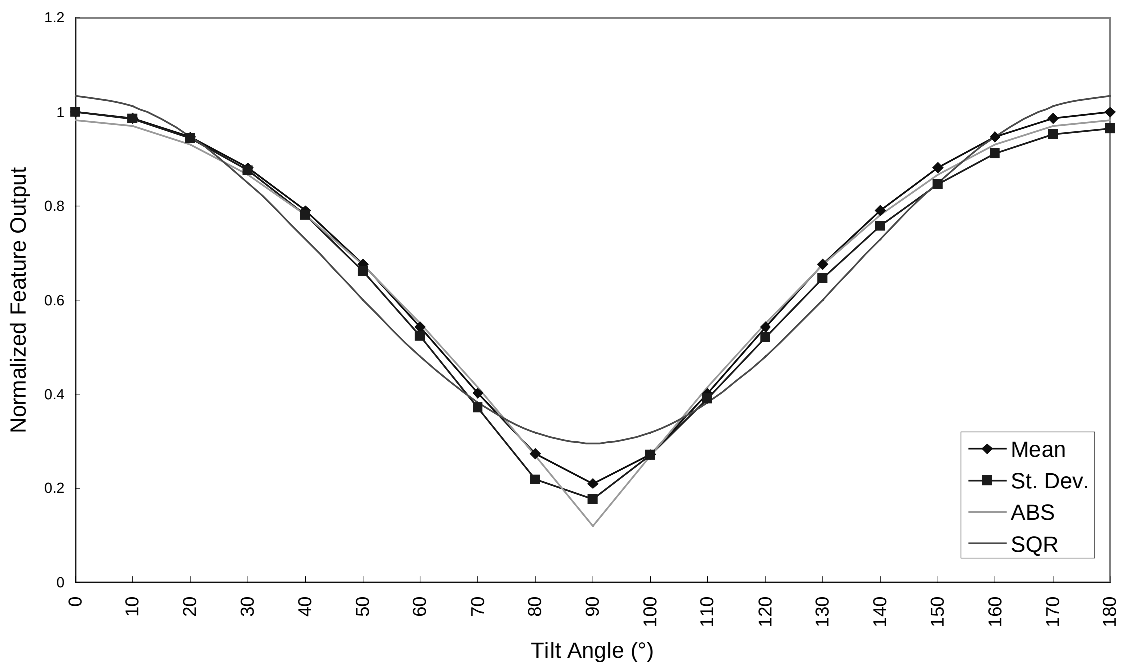

Figure 6.3.1 The effect of tilt angle variation on feature output

We note several points from the above graph:

(i) the mean and standard deviation have very similar tilt responses,

(ii) both have a similar a/b ratio,

(iii) both the mean and the standard deviation appear to be most successfully modelled using the absolute function. This is in agreement with the theoretical predictions.

In this thesis we will adopt the $a|\cos(\phi-\tau)|+b$ relationship as an approximation for the dependency of both the mean and standard deviation of the feature images on the illuminant tilt.

In the previous graph we only considered the case of a single filter, in any application it would be normal to use a larger feature set. In *Figure 6.3.2* we apply all the filters of the set within a frequency band, each is scaled by the same amount. As was predicted in Eq. 6.2.2e, the waveforms are effectively phase shifted versions of that shown in *Figure 6.3.1* each of which reaches its maximum when the tilt angle coincides with the filter orientation, i.e. $\tau=\phi$.

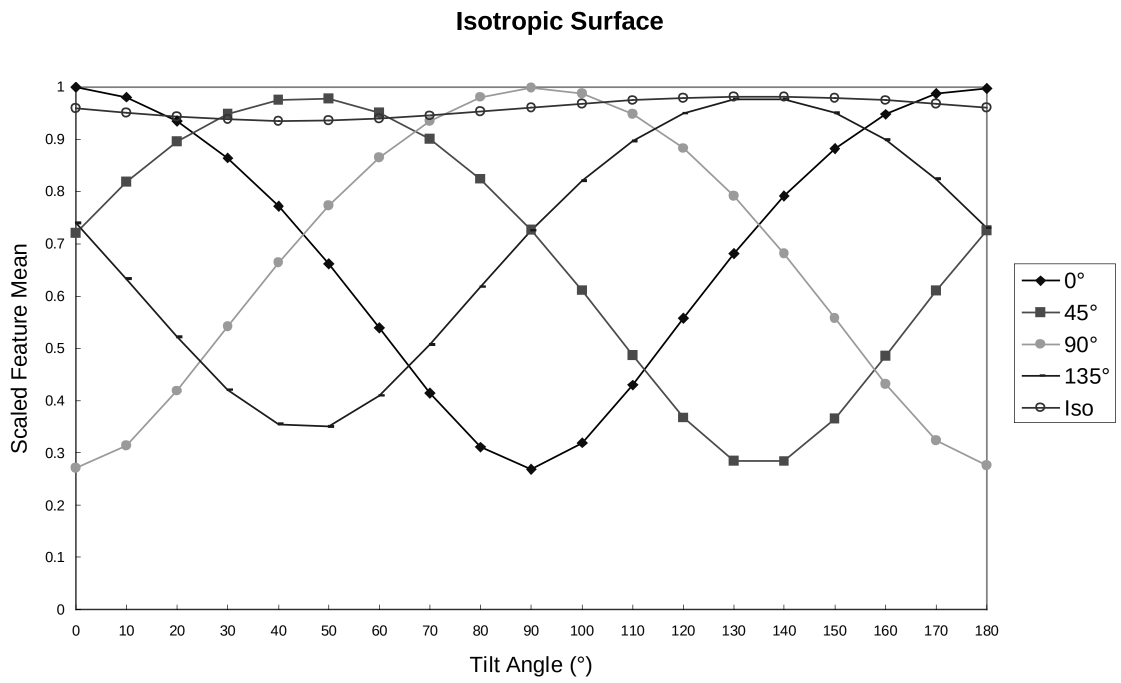

Figure 6.3.2 Tilt effects on the means of features obtained over a range of orientations in the F25 band.

Since the waveforms are phase shifted versions of each other, varying the illuminant tilt may be considered as being equivalent to rotation of the image in the case of an isotropic texture. To demonstrate this point, we plot the mean output of a fifth feature. This feature is derived from an isotropic, Gaussian bandpass filter[1]. As expected, this is almost unaffected by tilt variation.

We now consider the effect on feature output of tilt variation on the directional *Ogilvy* surface, again the feature means have been scaled by the maximum value of the largest waveform.

[1] Developed by T. Wittig

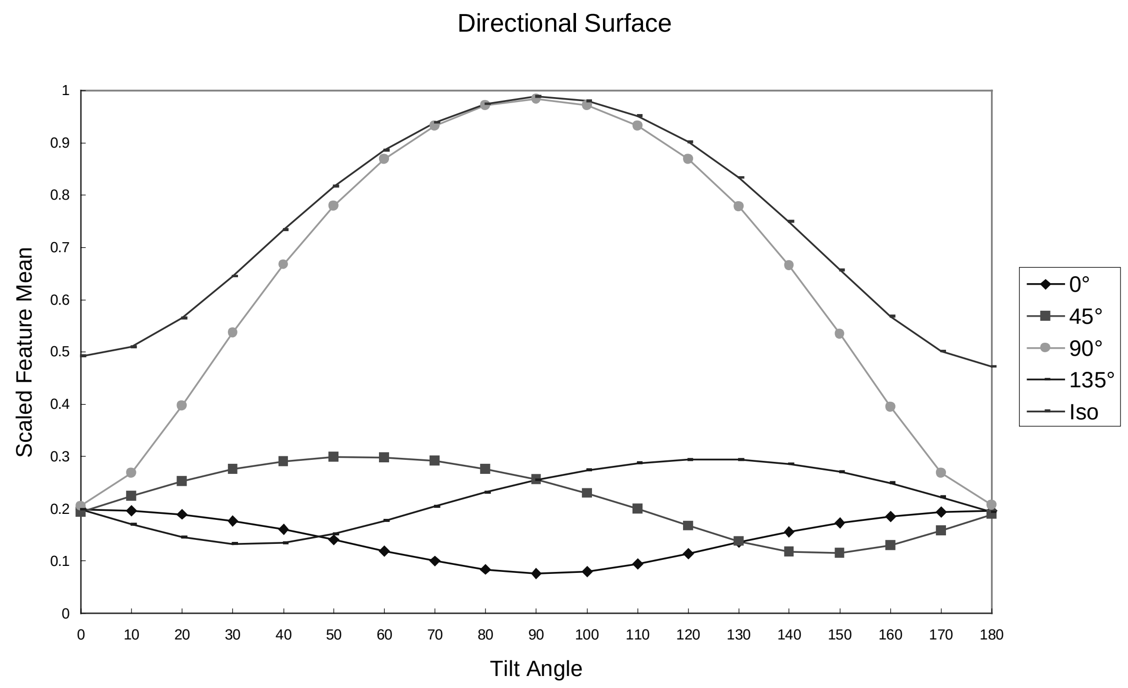

Figure 6.3.3 F25 Filter responses to directional texture with varying tilt.

We note the following points:

(i) on visual inspection the $|\cos(\phi-\tau)|$ relationship appears to hold for both textures and all filters.

(ii) the output of the filter (ϕ=0°) orientated at right angles to the grain of the surface is much larger than that of the other filters in the band.

(iii) the feature mean of the isotropic filter is heavily dependent on the tilt direction.

6.3.2 Verification by Experiment

We now consider the two real test samples *Rock* and *Striate*. As with the isotropic synthetic surface, the feature means for the *Rock* texture all vary in a cosine manner of approximately equal amplitude, though with a slightly lower a/b ratio.

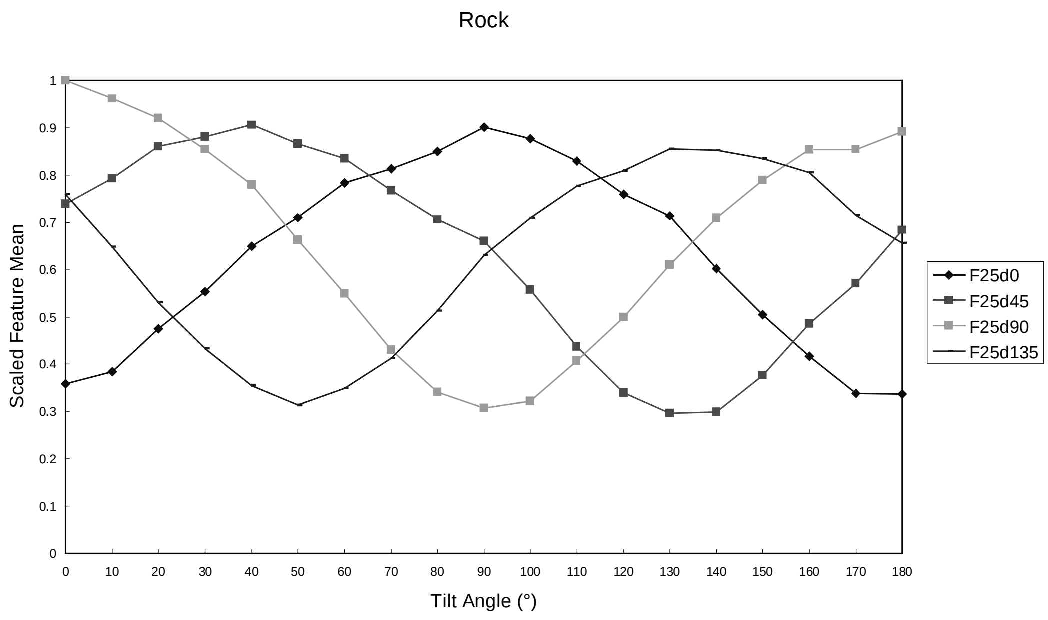

Figure 6.3.4 Filter responses to 'Rock' surface illuminated from various tilt angles.

The 'Striate' texture is analogous to the Ogilvy surface. As with the synthetic case all features vary in an approximately sinusoidal manner. The filter means can all be approximated by sinusoids of the same amplitude save that feature measured perpendicularly to the grain, which has a much larger amplitude.

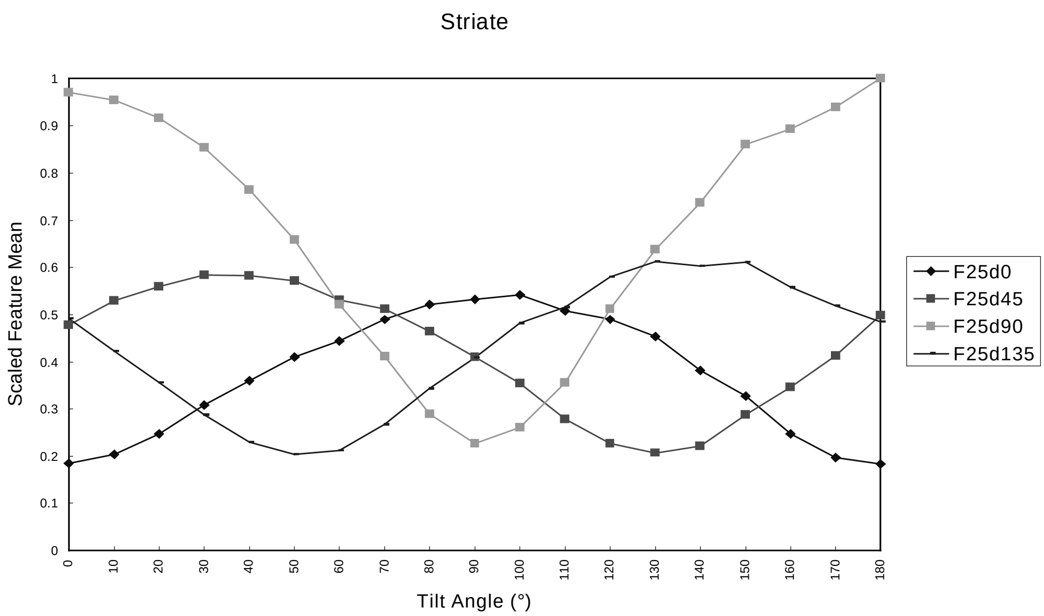

Figure 6.3.5 Filter responses to 'Striate' texture.

Although the feature means are important, in the context of discrimination the standard deviation of the feature must also be considered. In *Figure 6.3.6* and *Figure 6.3.7* the means and standard deviations of features derived from filters oriented at ϕ=0° and 90° are plotted for the *Rock* and *Striate* surfaces respectively. An $a\left|\cos(\phi-\tau)\right|+b$ curve is also shown in each graph where the *a* and *b* parameters have been estimated using least squares.

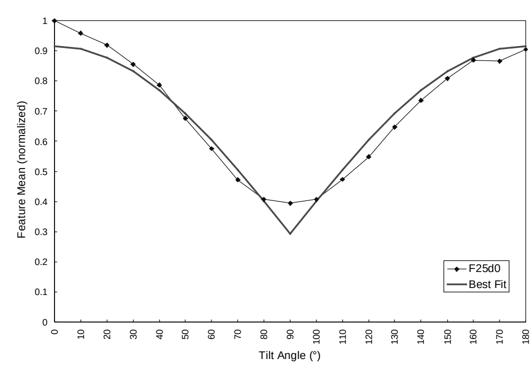

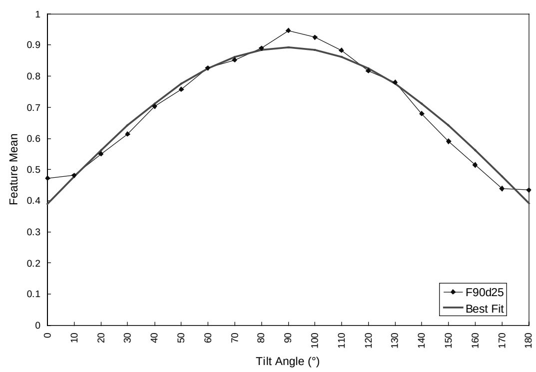

F25d0 Mean

F25d90 Mean

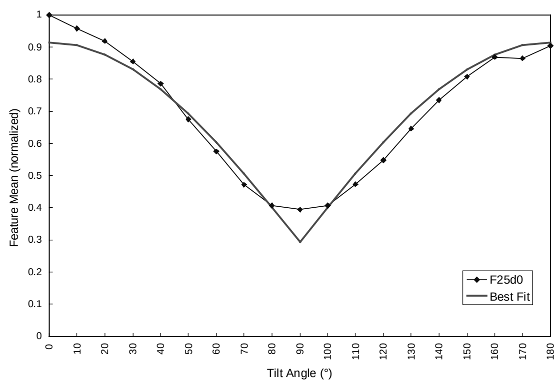

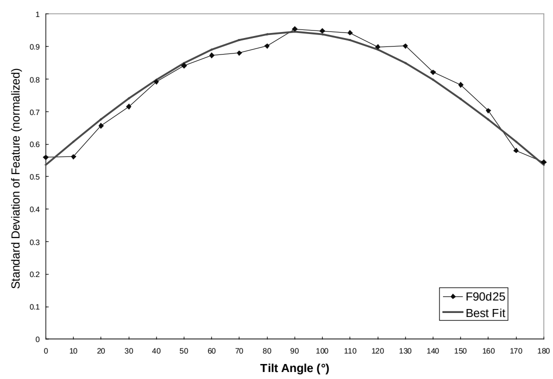

F25d0 Standard deviation

F25d90 Standard deviation

Figure 6.3.6 Variation of Rock feature statistic with tilt angle.

In the case of *Rock* the feature mean fits well to the parameterised curve, though the minima of the F25d0 graphs are much smoother than predicted. Both the mean and standard deviation of the F25d0 features derived from *Striate* appear to be a better fit in this respect, though there is a small phase disparity between the predicted and the observed curves in the F25d90 case. Nevertheless, the proposed relationship does appear to be a reasonable approximation to the experimental findings.

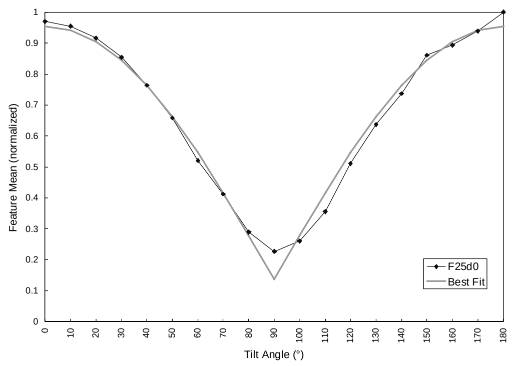

(a.) F25d0 Mean

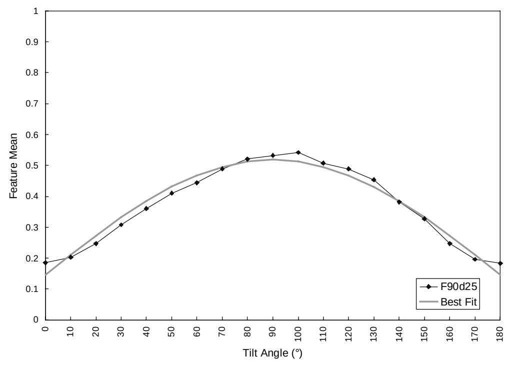

F25d90 Mean

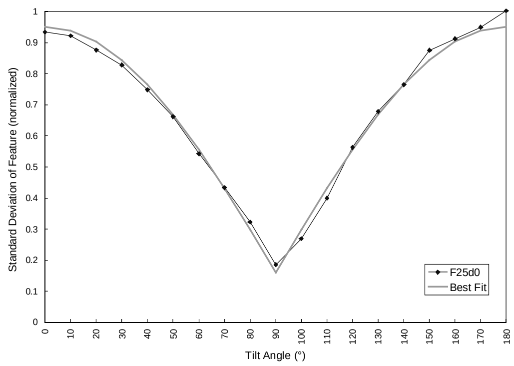

F25d0 Standard deviation

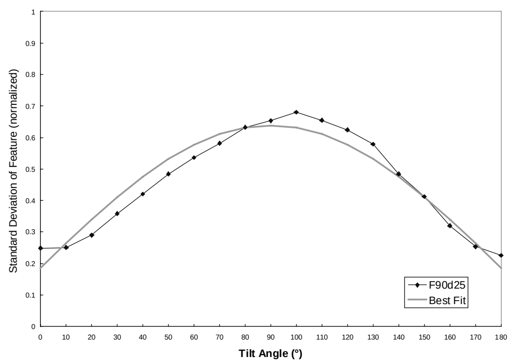

F25d90 Standard deviation

Figure 6.3.7 Variation of Striate feature statistics with tilt angle.

In *Table 6.3.1* and *Table 6.3.2* the estimated parameters for the feature mean/tilt relationship are tabulated for the *Rock* and *Striate* surfaces respectively. It was noted earlier that the mean curve of the F25d0 feature for the *Striate* texture is much larger than the means of the other features in the set. The *a* (tilt-dependent) parameter reflects this; the F25d0 feature parameter is almost twice that of the other features. The *b* parameter estimate for F25d0 is noticeably larger than for the features, however, the difference is much less significant than was observed with the *a* parameter.

Striate	Mean			Standard Deviation		
Filter	a	b	a/b	a	b	a/b
F25d0	14.077	5.987	2.35	2.440	1.181	1.348
F25d45	7.702	4.740	1.62	1.608	0.969	1.659
F25d90	6.870	4.125	1.67	1.534	0.938	1.636
F25d135	7.086	4.914	1.44	1.839	1.083	1.698

Table 6.3.1 Feature statistics for Striate surface.

The parameter estimates for the *Rock* show a significant amount of variation in the *b*-parameter, although the variation is small relative to that observed in the *Striate* parameters. The *a* parameter is more consistent for the *Rock* than for the *Striate* surface, and the variation which does exist does not show the same degree of linkage to the *b*-parameter as in the *Striate* estimates.

Rock	Mean			Standard Deviation		
Filter	a	b	a/b	a	b	a/b
F25d0	9.245	5.642	1.64	1.522	1.043	1.46
F25d45	7.606	5.687	1.34	1.321	1.147	1.15
F25d90	7.950	5.642	1.41	1.498	1.235	1.21
F25d135	8.736	5.205	1.68	1.101	1.162	0.947

Table 6.3.2 Feature statistics for Rock surface.

The results obtained are to some degree ambiguous, however, we make the following conclusions, albeit based on a very limited data set: the *a* parameter is largely dependent on the nature of the surface in the direction of the filter—it will therefore vary widely from surface to surface. The *b*-parameter also contains an element of dependency on the surface characteristics within the filter's bandpass region, however, it is largely invariant to the surface characteristics.

6.3.3 Summary of Classifier Modelling

In this section we have shown that measures derived from directional Gabor filters vary with tilt. Both mean and standard deviation vary with a relationship approximated by Eq. 6.3.3.

$$a\left|\cos(\phi - \tau)\right| + b \tag{6.3.3a}$$

We have shown this using theory, simulation and experiment.

We have also shown that an isotropic filter, applied to the image of an isotropic surface is unaffected by tilt variations. However, an isotropic filter applied to the image of a directional texture *will* be affected by tilt variation. In addition, an isotropic feature will not, of course be able to capture important directional information.

By demonstrating that tilt affects the output of features, we have shown that there exists a potential problem for classification. In the following sections we show tilt induced classification failure occurring for several texture classification tasks.

6.4 The Effect on a Classifier

We shall now consider the effect of tilt variation in terms of classification accuracy. Due to interdependency of classes associated with classification, any treatment of this subject is necessarily empirical. We use the controllability of the synthetic textures to define a relatively simple classification task. This enables us to reduce the feature set to just two features (Gabor filters orientated to 0° and 90°) and still maintain a good level of classification. With only two features we may characterise the classification process with a two dimensional scatter plot. We also plot the feature measures for each texture in *Figure 6.4.1*. This section is designed to use this window into classification to gain an intuitive understanding of how clusters behave during tilt variation and how this will affect classification.

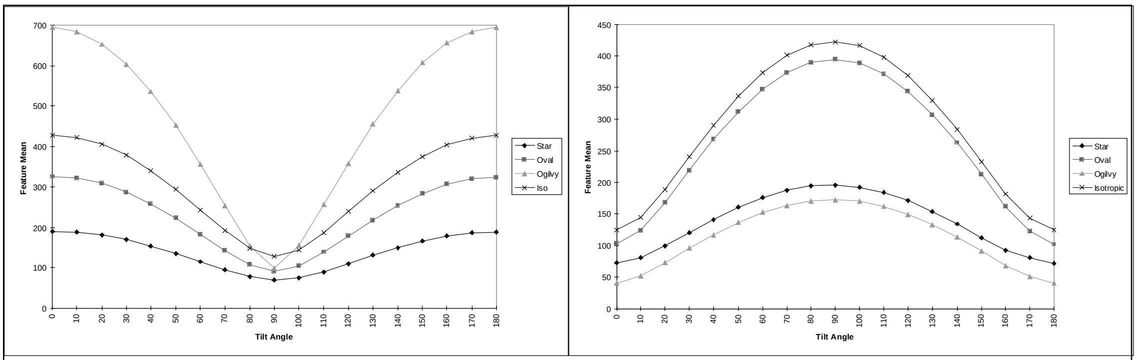

Figure 6.4.1 Feature Means derived from F64d0 and F64d90 Filters

The classifier is trained on surfaces illuminated from τ=0°, classification at this stage being generally good, the most prominent misclassification being mutual confusion between the isotropic and mildly directional surfaces, (see *Figure 6.4.2)*. The relatively compact nature of the clusters in the vertical direction reflects the attenuation of the vertical frequencies.

As the illumination is rotated towards the vertical, the vertical frequencies are accentuated and the horizontal frequencies attenuated. The cluster centres now begin to trace out an approximation to simple harmonic motion in feature space whereas the clusters themselves contract and expand along their feature axis. By the time the illuminant has reached the vertical, all the clusters lie in the area of feature space assigned to the surface which displayed most vertical energy at τ=0°, which in this case is the isotropic surface. Misclassification of the other textures is almost complete.

Figure 6.4.8

6.5 Classification Experiments For Real Textures

The thesis of this chapter is that variation in illuminant tilt between training and classification can cause the classifier to fail catastrophically. In this section, this effect will be demonstrated experimentally using real data.

6.5.1 Test Criteria

In order to demonstrate this effect, we must show three things:

1. The classifier performs well when the tilt angle is identical for training and classification.
2. The misclassification rate increases progressively, though not necessarily linearly, with the cosine of the angle between the illuminant vectors of the training and classification,
3. Implied from (2) is that the misclassification rate at τ=180° should be approximately equal to that at τ=0°.

6.5.2 The Data Set

In this section we will use three texture montages: *Anaglypta, Stones1* and *Stones2,* shown in *Figure 6.5.1*. The *Anaglypta* montage consists of highly directional, and highly uniform textured surfaces. It consequently represents the easiest classification task. The *Stones1* montage consists of three approximately isotropic rock surfaces and one highly directional surface, *Striate,* in which the directionality is aligned approximately with the Y-axis, (θ=0°). The textures comprising the *Stone2* montage are all directional to some degree. The directionality of the *Slate* and *Pitted* surfaces is aligned with the θ=0° direction. The *Twins* direction has its most prominent directionality in the direction θ=90°, while the radial texture is directional at θ=45°.

Anaglypta Montage

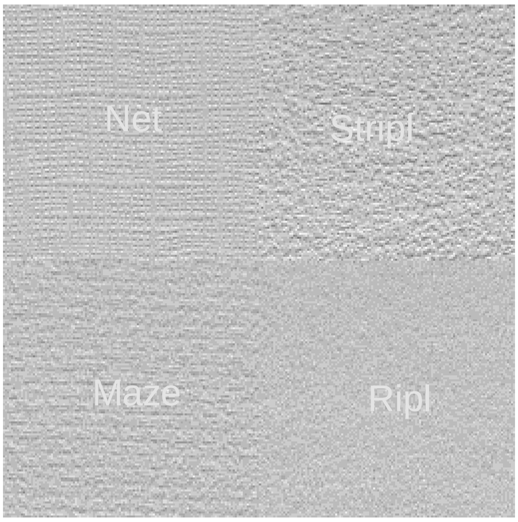

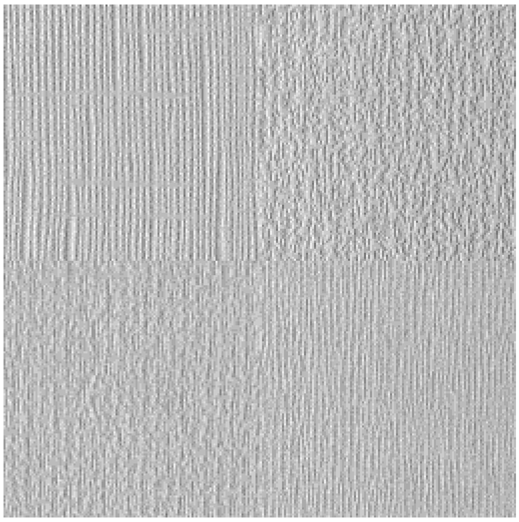

Stone 1 Montage

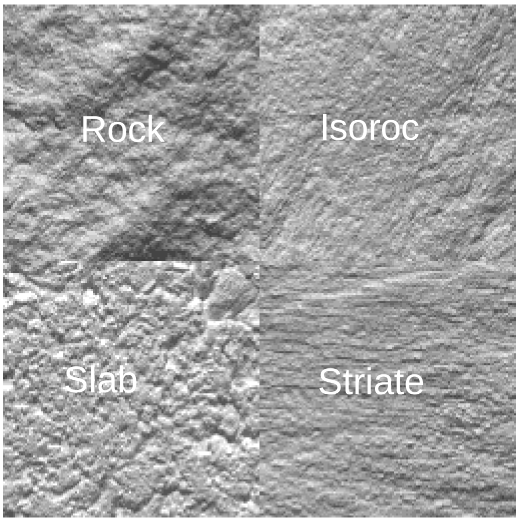

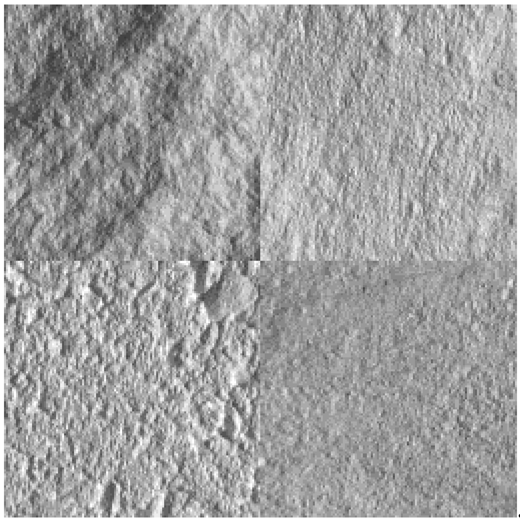

Stone 2 Montage

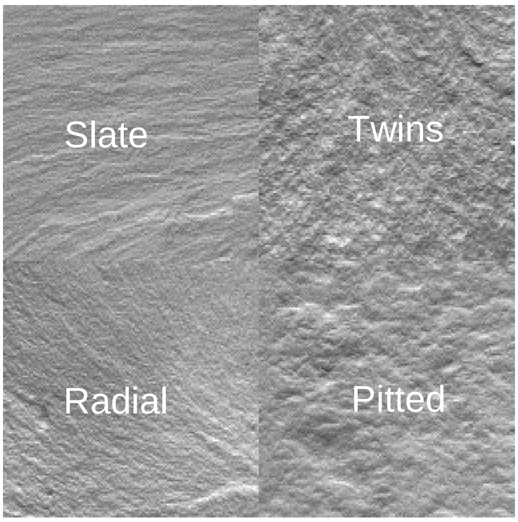

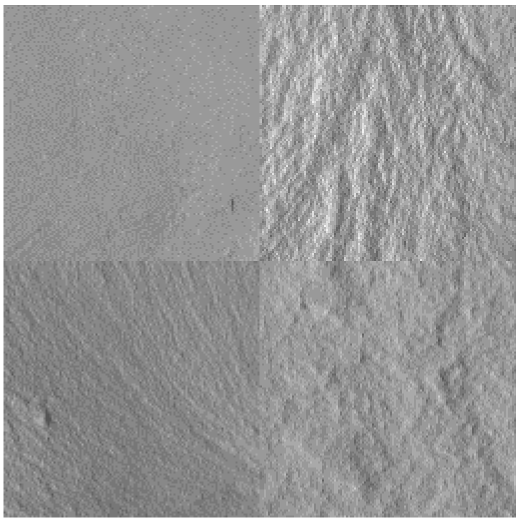

Figure 6.5.1 Test montages of surfaces, illuminated at Tau 0° (left) and 90° (right).

6.5.3 Experimental Work

First Criterion

Our first criterion is that the classifier should be perform classification accurately for the τ=0° image. Real textures were classified using the same system as the synthetics in the previous section, though a much larger feature set with twelve members and a 12*12 mode postprocessing filter were used to achieve an acceptable classification accuracy. Unfortunately the ability of the classifier to cope with the heterogeneous nature of surface roughness on the samples is poor, and to obtain an acceptable level of accuracy, it was necessary to use the entire montage for training, this means that we are only testing the ability of the features to describe the textures and are not testing the classifiers ability to generalise from the training data. The classified images and the misclassification rates are shown in *Figure 6.4.2*.

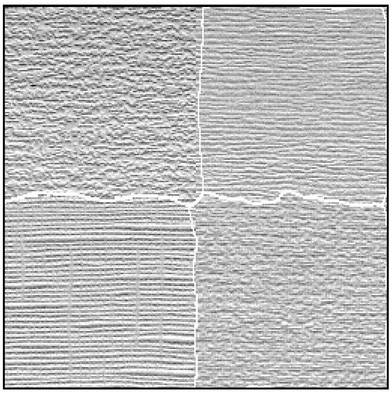

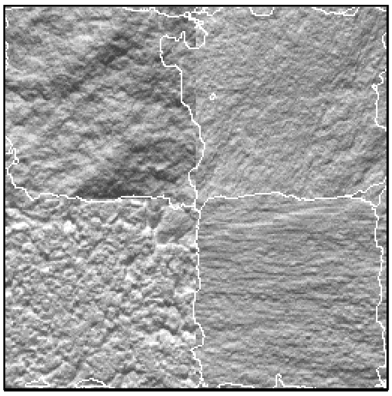

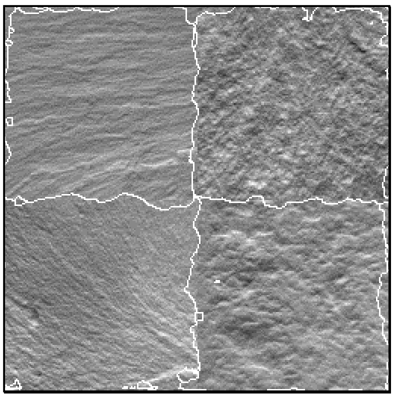

Montage	Misclassification (%)
Anaglypta	2.88
Stones1	5.09
Stones2	3.29

Figure 6.5.2 Montages trained and classified at Tau 0°.

As predicted, the anaglypta montage is the easiest classification. Whereas the Stones montages are classified to a much lower degree of accuracy.

Second Criterion

The second criterion for our thesis is that the classifier should be progressively less accurate as tilt angle increasingly differs from the training angle. The classifier was trained at τ=0° and then tested on images of the surfaces illuminated from τ=0 to 180° in 10° degree increments. The degree of misclassification for this classifier was recorded for each illumination condition in *Figure 6.5.3-5* In addition to the classifier trained at 0°, a classifier retrained for each tilt angle, labeled "*Best*", was also employed. This acts as a control in displaying the level of difficulty of classification inherent in a particular classification, which enables us to resolve the misclassification rate which is due to inappropriate training data.

In only one of the three cases does the misclassification rate increase in an approximately monotonic fashion as the cosine of the tilt angle increases—fulfilling our second criterion. However, deviations from the expected behaviour occur in regions of high misclassification, and are still of a magintude such that classfication at these tilt angles, with this classifier, is pointless. In this way we argue that the second criterion has been fulfilled to a satisfactory degree.

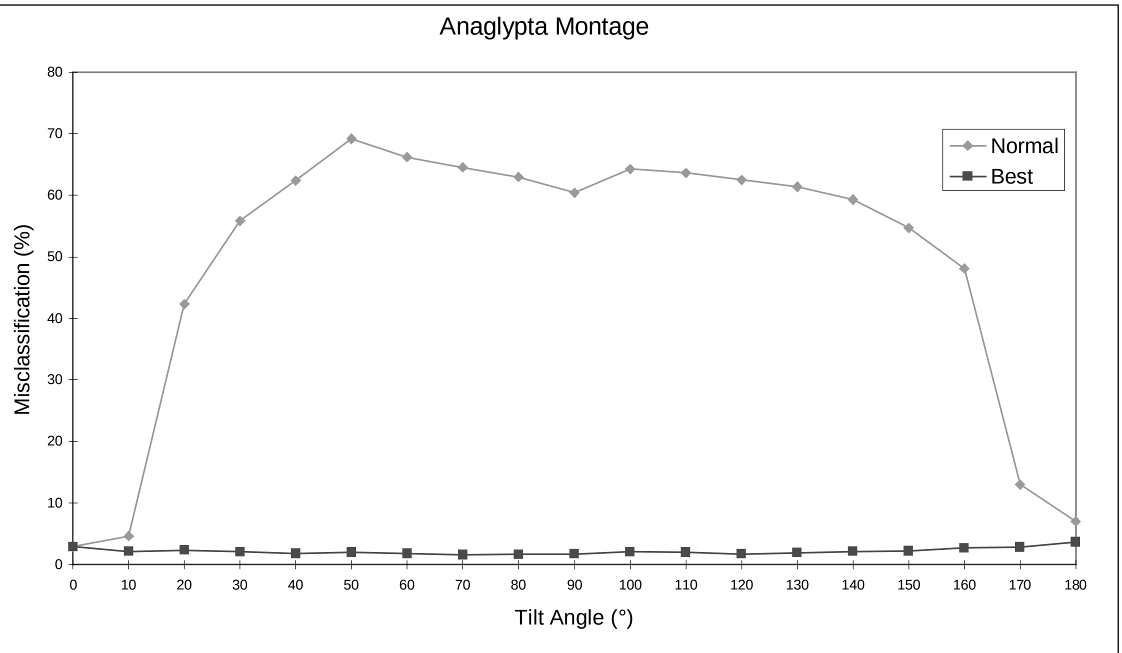

Figure 6.5.3 Misclassification rates for Anaglypta montage.

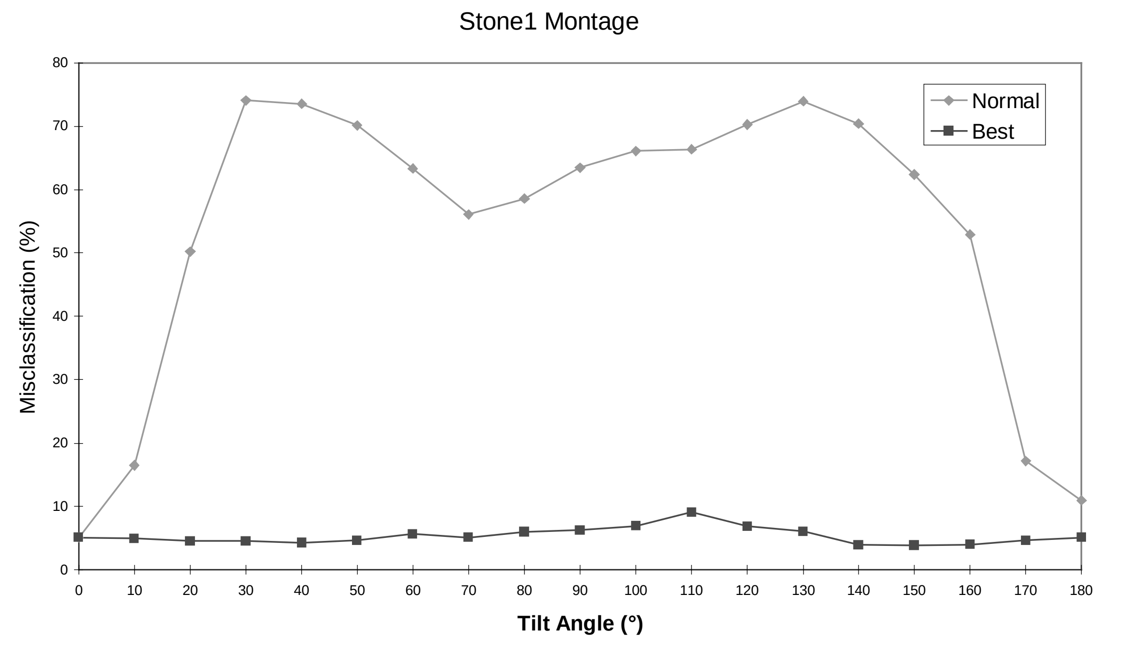

Figure 6.5.4 Misclassification rates for Stones 1 montage.

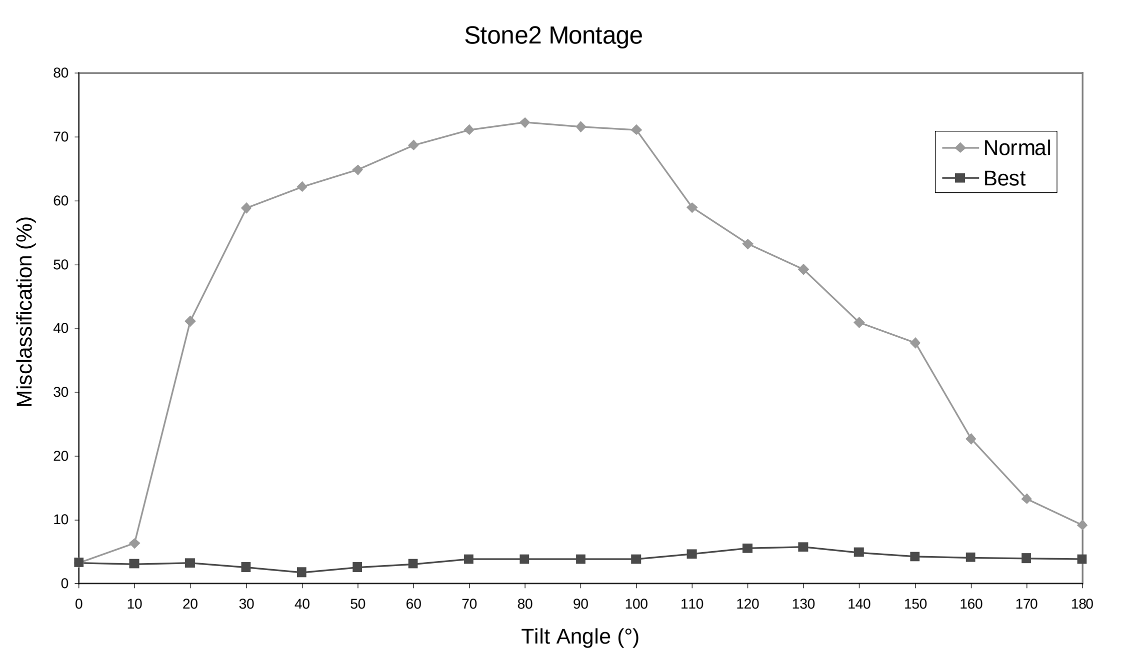

Figure 6.5.5 Misclassification rates for Stone2 montage.

Third Criterion

In order to show that the increased rate of misclassification is due to the effect of illuminant tilt our third criterion, i.e. an accurate classification at τ=180°, must be met. In all cases the classification at τ=180° is significantly poorer than that at 0°, however, in the context of the classifications at intermediate values of τ, we believe the third criterion has been fulfilled .

	Misclassification at τ=0°	Misclassification at τ=180°
Anaglypta	2.88	4.30
Stone1	5.09	10.95
Stone2	3.29	9.19

Table 6.5.1 Comparison of misclassification rates.

6.5.4 Summary of results

In order to show that classification is dependent on the illuminant tilt angle we set three criteria:

1. the classifier must be able to classify surfaces imaged under the same illuminantion conditions as those at which the training data was obtained,
2. the level of misclassification should increase with the cosine of the difference between the tilt angles at training and classification,
3. the classifier should be able to classify the τ=180° image accurately,

Using experiments on real textures, it was shown that these criteria were fulfilled to a level that provides strong evidence for the tilt dependency. We therefore argue that classification is tilt dependent, and, where illuminant tilt cannot be held constant between training and classification, the naive classifier developed in the previous chapter is not adequate.

6.6 Summary

At the beginning of this chapter we stated two aims: to model the effect of illuminant tilt on the classifier, and to observe the effect of tilt on the accuracy of classification.

Sections 6.2 and 6.3 were concerned with modelling the tilt response of the features. In the first section, the analytical model of the imaging process was extended to include the linear stage of the classifier, i.e. Gabor filtering. The analysis allowed the prediction that the first order statistics of the feature images would vary with an $a|\cos(\phi-\tau)| + b$ relationship. This was verified approximately by the simulations and experiments performed in section 6.3.

While the calculation of features is amenable to analysis, the process of discrimination is inherently non-linear and cannot be integrated into our model. In order to extend our analysis we used a synthetic classification task to observe the effect of tilt on the movement of feature clusters across discriminant boundaries. Finally the degradation in classification accuracy due to tilt variation was shown on three montages of real textures.

In this chapter, we have shown, using theory, simulation and experiment, that varying the tilt angle of the illuminant induces the movement of feature clusters. Where training and classification images are obtained at different tilt angles, this movement may cause clusters to move across discriminant boundaries. Using simulation and experiment, it was shown that this may occur to such a degree as to seriously degrade the performance of the classifier. In the next chapter we will consider several schemes to mitigate this effect.