The Effect of Illuminant Rotation on Texture Filters: Lissajous's Ellipses

M. Chantler[1], M. Schmidt[2], M. Petrou[3], and G. McGunnigle[1]

[1] Texture Lab., Heriot-Watt University, Edinburgh,Scotland,
mjc@cee.hw.ac.uk
http://www.cee.hw.ac.uk/texturelab/
[2] Darmstadt University of Technology, Germany,
[3] Department of Electronic and Electrical Engineering, University of Surrey, Guildford,
Surrey, GU2 7XH, United Kingdom
M.Petrou@ee.surrey.ac.uk
http://www.ee.surrey.ac.uk/Personal/M.Petrou

Abstract. Changes in the angle of illumination incident upon a 3D surface texture can significantly change its appearance. These changes can affect the output of texture features to such an extent that they cause complete misclassification. We present new theory and experimental results that show that changes in illumination tilt angle cause texture clusters to describe Lissajous's ellipses in feature space. We focus on texture features that may be modelled as a linear filter followed by an energy estimation process e.g. Laws filters, Gabor filters, ring and wedge filters. This general texture filter model is combined with a linear approximation of Lambert's cosine law to predict that the outputs of these filters are sinusoidal functions of illuminant tilt. Experimentation with 30 real textures verifies this proposal. Furthermore we use these results to show that the clusters of distinct textures describe different elliptical paths in feature space as illuminant tilt varies. These results have significant implications for illuminant tilt invariant texture classification.

Keywords: Texture, illumination, texture features

1 Introduction

It has been shown that changes in the angle of illumination incident upon a 3D surface texture can change its appearance significantly as illustrated in Fig. 1. Such changes in image texture can cause complete misclassification of surface textures [1]. Essentially the problem is that side-lighting, as used for instance in Brodatz's texture album [2], enhances the appearance of surface texture but produces an image which is a directionally filtered version of the surface height function. Furthermore as the theory developed by Kube and Pentland [3] predicts, the axis of this filter is a function of the illumination's tilt angle [1] . This is

[1] In the axis system we use, the camera axis is parallel to the z-axis, illuminant tilt is the angle the illuminant vector makes with the x-axis when it is projected into the

A. Heyden et al. (Eds.): ECCV 2002, LNCS 2352, pp. 289–303, 2002.

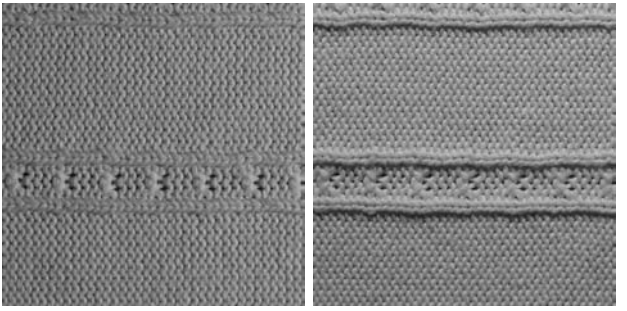

Fig. 1. Two images of the same surface texture sample captured using different illuminant tilt angles

unfortunate as many texture classifiers employ directional filters in their feature measures [4]. It is not surprising therefore that rotation of the directional filter formed by the illumination function can cause conventional image-based texture classifiers to fail dramatically. Hence the objective of this paper is to present theory and experimental results that allow the behaviour of texture features in classification space to be predicted as a function of the illumination's angle of tilt. This provides both an insight into the behaviour of existing classifiers and a basis for the development of tilt invariant classifiers.

Very little work has been published on this subject with the exception of [5] which presents an informal proof that texture feature means are sinusoidal functions of the illumination's tilt angle. Dana, Nayer, van Ginneken and Koenderink established the Columbia-Utrecht database of real world surface textures which they used to investigate bidirectional texture functions [6]. Later they developed histogram [7,8] and correlation models [9] of these textures. Leung and Malik developed a texture classification scheme that identifies 3D 'textons' in the Columbia-Utrecht database for the purposes of illumination and viewpoint invariant classification [10,11]. Neither of these papers is concerned with developing a theory for the tilt response of texture features. In this paper we:

1. present formal theory for the behaviour of texture feature means that also provides expressions for the coefficients of these sinusoidal functions;
2. provide experimental data from thirty real textures [2] that support this theory;
3. use this theory to show that the behaviours of texture clusters follow super-elliptical trajectories in multi-dimensional feature spaces; and finally

x, y plane, and illuminant slant is the angle that the illuminant vector makes with the camera axis.

[2] We have not used the Columbia-Utrecht database as the illumination was held constant while the viewpoint and orientation of the samples were varied during data capture.

4. show experimental results in 2D feature space that support the conjecture that changes in illumination tilt cause texture clusters to follow behaviours that may be described by Lissajous's ellipses.

We focus on the group of texture features classified as 'filtering approaches' by Randen [4] such as Gabor, Laws, and ring/wedge filters. These features may be described as a linear filter followed by an energy estimation function. This general feature model is combined with a linearised version of Lambert's cosine law due to Kube and Pentland [3]. The result is theory that may be used to predict the behaviour of texture features as a function of illuminant tilt. This theory is presented in the next section, after which we describe results derived from experiments with 30 textures.

2 Theory

This section derives an expression for the mean value of a texture filter as a function of the illumination's tilt angle. The features are modelled simply as a linear filter followed by an energy estimation process as described in the introduction. First, however, we give a short derivation of the linear illumination model that is necessary for the development of the theory.

2.1 A Linearised Model of Lambert's Cosine Law

We assume that the surface has Lambertian reflection, low slope angles, and that there is no significant shadowing or inter-reflection. Ignoring the albedo factor Lambert's cosine rule may be expressed as:

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

where:

$i(x,y)$	is the radiant intensity;
$p(x,y)$	is the partial derivative of the surface height function in x direction;
$q(x,y)$	is the partial derivative of the surface height function in y direction;
τ	is the tilt angle of the illumination; and
σ	is the slant angle of the illumination respectively.

For p and $q \ll 1$ we can use a truncated Taylor's series to linearise this equation to:

$$i(x,y) = -\cos(\tau)\sin(\sigma)p(x,y) - \sin(\tau)\sin(\sigma)q(x,y) + \cos(\sigma)$$

Transforming the above into the frequency domain and discarding the constant term we obtain:

$$\begin{aligned} & \mathcal{I}(\omega,\theta) = [-\cos(\tau)\sin(\sigma)\, i\,\omega\, \cos(\theta) - \sin(\tau)\sin(\sigma)\, i\,\omega\, \sin(\theta)]\mathcal{H}(\omega,\theta) \\ \iff & \mathcal{I}(\omega,\theta) = -i\,\omega\, \sin(\sigma)\cos(\theta-\tau)\mathcal{H}(\omega,\theta) \end{aligned} \tag{2}$$

where:

$\mathcal{I}(\omega, \theta)$ is the Fourier transform of the image intensity function;
$\mathcal{H}(\omega, \theta)$ is the Fourier transform of the surface height function; and
(ω, θ) are polar frequency coordinates.

In this paper it is more convenient to express equation 2 in its power spectrum form:

$$I(\omega, \theta) = \omega^2 cos^2(\theta - \tau) sin^2(\sigma) H(\omega, \theta) \tag{3}$$

where:

$I(\omega, \theta)$ is the image power spectrum; and
$H(\omega, \theta)$ is the surface power spectrum.

Equations 2 and 3 are similar to the model first expressed by Kube & Pentland in [3]. In the context of this paper the most important feature of this model is the $\cos(\theta - \tau)$ factor, which predicts that the imaging function acts as a directional filter of the surface height function.

2.2 The Output of Linear Texture Filters and Their Features

We define a Linear Texture Feature as a linear filter followed by a variance estimator[4]. The process formed by applying such a feature to an image is as shown in Fig. 2. Since the model of the illumination process (equation 3) is also linear we may exchange it with the linear filter (Fig. 3). We use $A(\omega, \theta)$ to represent the notional power spectrum of the output of the linear texture filter applied directly to the surface height function.

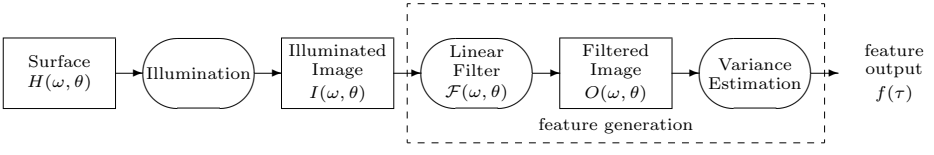

Fig. 2. Feature generation

Fig. 3. The feature generation model with illumination and linear filter processes interchanged

Thus to determine the mean output of a Linear Texture Filter we simply have to apply Kube's model in the form of equation 3 to $A(\omega, \theta)$ and develop an expression for the variance of the subsequent output.

The mean output of a linear texture feature is the variance of the output $o(x, y)$ of its linear filter:

$$f(\tau) = \mathcal{VAR}(o(x, y)) \tag{4}$$

If we assume that $o(x, y)$ has a zero mean and that $O(\omega, \theta)$ is its power spectrum expressed in polar co-ordinates then we may express equation 4 as:

$$\begin{aligned} f(\tau) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} O(\omega, \theta)\, du\, dv \\ \iff f(\tau) &= \int_{0}^{\infty}\int_{0}^{2\pi} \omega\, O(\omega, \theta)\, d\theta\, d\omega \end{aligned} \tag{5}$$

Using equation 3 we can express $O(\omega, \theta)$ as follows:

$$\begin{aligned} O(\omega, \theta) &= |\mathcal{F}(\omega, \theta)|^2 \omega^2 cos^2(\theta - \tau) sin^2(\sigma) H(\omega, \theta) \\ \iff O(\omega, \theta) &= \omega^2 cos^2(\theta - \tau) sin^2(\sigma) A(\omega, \theta) \end{aligned} \tag{6}$$

where:

$\mathcal{F}(\omega, \theta)$ is the transfer function of the linear filter; and
$A(\omega, \theta)$ $= H(\omega, \theta)|\mathcal{F}(\omega, \theta)|^2$

Substituting equation 6 into equation 5 we obtain:

$$f(\tau) = \int_{-\infty}^{\infty} \omega^3 sin^2(\sigma) \int_{0}^{2\pi} cos^2(\theta - \tau) A(\omega, \theta)\, d\theta\, d\omega \tag{7}$$

In a previous paper [5] we used an informal argument to show equation 7 is a sinusoidal function of τ. However, in this paper we make a simple trigonometrical substitution.

Using:

$$cos^2(x) = 1/2\,(1 + cos(2x)) \text{ and } cos(x - y) = cos(x)cos(y) + sin(x)sin(y)$$

gives:

$$\begin{aligned} f(\tau) &= \int_{0}^{\infty} \omega^3 sin^2(\sigma) \int_{0}^{2\pi} 1/2\,[1 + cos(2\theta)cos(2\tau) + sin(2\theta)sin(2\tau)] A(\omega, \theta)\, d\theta\, d\omega \\ &= a + b\cos(2\tau) + c\sin(2\tau) \\ &= a + d\cos(2\tau + \phi) \end{aligned} \tag{8}$$

where:

$$a = 1/2\, sin^2(\sigma) \int_0^{\infty} \omega^3 \int_0^{2\pi} A(\omega, \theta)\, d\theta\, d\omega$$

$$b = 1/2\, sin^2(\sigma) \int_0^{\infty} \omega^3 \int_0^{2\pi} cos(2\theta) A(\omega, \theta)\, d\theta\, d\omega$$

$$c = 1/2\, sin^2(\sigma) \int_0^{\infty} \omega^3 \int_0^{2\pi} sin(2\theta) A(\omega, \theta)\, d\theta\, d\omega$$

$$d = \sqrt{b^2 + c^2}, \qquad \phi = \arctan(c/b)$$

The above parameters (a, b etc.) are all functions of illuminant slant (σ) and $A(\omega, \theta)$, which is itself a function of the surface height function and the linear texture filter. None of them is a function of illuminant tilt (τ). Thus equation 8 predicts that the output of a texture feature based on a linear filter is a sinusoidal function of illuminant tilt [3] with a period π radians.

2.3 Behaviour in a Multi-dimensional Feature Space

As texture classifiers normally exploit the output of several features it is important to investigate the behaviour of texture features in multi-dimensional decision space.

If two different features are derived from the same surface texture, the results can be plotted in a two-dimensional x, y feature space. Using our sinusoidal prediction we obtain the general behaviour:

$$x = f_1(\tau) = a_1 + b_1 \cos(2\tau + \phi_1)$$

$$y = f_2(\tau) = a_2 + b_2 \cos(2\tau + \phi_2)$$

Since the frequency of the two cosines is the same these two equations form two simple harmonic motion components. Therefore the trajectory in 2D feature space is in general a Lissajous ellipse.

There are two special cases. If the surface is isotropic and the two filters are identical except for a difference in direction of 90°, the mean value and the oscillation amplitude of the two features are the same and the phase difference becomes 180°. Thus we predict that the scatter plot for isotropic textures and two identical but orthogonal filters is a straight line.

[3] In the case of $A(\omega, \theta)$ being isotropic (for instance if both the surface and the filter are isotropic) the response will degenerate to a sinusoid of zero amplitude, i.e. it will be a constant (straight-line) function of τ. However, if an isotropic filter is applied to a directional surface then $A(\omega, \theta)$ will not be isotropic and the tilt response will be a sinusoidal function of tilt.

If the surface is isotropic and the two filters are identical except a difference in direction of 45°, the mean value and the oscillation amplitude of the two features are again the same but the phase difference is now 90°. In this case we predict that the scatter plot is a circle.

The line and the circle are the two special cases of all possible curves. In the general case of two or more filters the result is an ellipse or super-ellipse.

2.4 Summary of Theory

In this section we derived a formula of the behaviour of texture features as a sinusoidal function of illuminant tilt. The parameters of this equation $f(\tau) = a + d\,cos(2\tau + \phi)$ are dependent on the height map of the surface and the linear texture filter. We also predicted that the resulting figure in a multi-dimensional feature space is a super-ellipse.

3 Assessing the Validity of the Sinusoidal Output Prediction

The theory in the previous section predicts that the outputs of texture features derived from linear filters are sinusoidal functions of the illuminant tilt angle and share a common form:

$$f(\tau) = a + d\cos(2\tau + \phi) \tag{9}$$

The purpose of this section is to present results from a set of experiments that were designed to assess the validity of this proposition. We used 30 different real surfaces and filtered them with six Gabor and two Laws filters. Two different slant angles were used and this resulted in a total of 324 datasets. Each dataset was assessed to see how closely it followed a sinusoidal function of 2τ using a goodness of fit error metric.

In this section we summarise the results using a histogram of the values of this error metric and present upper quartile, median and lower quartile results. First however, we will briefly describe the image-set, the texture features, and the error metric.

3.1 The Image-Set

Thirty physical texture samples were used in our experiments. 512x512 8-bit monochrome images were obtained from each sample using illumination tilt angles ranging between 0° and 180° degrees incremented by either 10° or 15° degree steps. All textures were illuminated at a slant angle of 45°. In addition six surfaces were also illuminated at a slant angle of 60°. The final dataset contains over 600 images. Table 1 contains one example image of every texture.

3.2 The Texture Features

Gabor filters [12] have been very popular in the literature so we have used six texture features of this type. In addition two Laws filters [13] were also used as they are extremely simple to implement, are very effective, and were one of the first sets of filters to be used for texture classification.

Gabor Filters. We use the notation typeFΩAΘ to denote a Gabor filter with a centre frequency of Ω cycles per image-width, a direction of Θ degrees, and of type complex or real. Five complex Gabor filters (comF25A0, comF25A45, comF25A90, comF25A135, comF50A45) together with one real Gabor filter (realF25A45) were implemented. All of the Gabor filters were implemented in the frequency domain.

Laws Filters. Laws [13] developed a set of two-dimensional FIR filters derived from three simple one-dimensional spatial domain masks.

L3 = (1,2,1) $\Longrightarrow$ "level detection"
E3 = (-1,0,1) $\Longrightarrow$ "edge detection"
S3 = (-1,2,-1) $\Longrightarrow$ "spot detection"

Laws reported that the most useful filters were a set of 5x5 filters which he obtained by convolving, and transposing his 3x1 masks. We have used two here: $L5E5$ and $E5L5$.

$$L5E5 = L5^T * E5 = (L3 * L3)^T * (L3 * E3) = \begin{bmatrix} -1 & -2 & 0 & 2 & 1 \\ -4 & -8 & 0 & 8 & 4 \\ -6 & -12 & 0 & 12 & 6 \\ -4 & -8 & 0 & 8 & 4 \\ -1 & -2 & 0 & 2 & 1 \end{bmatrix}$$

The $E5L5$ mask is simply the transpose of $L5E5$. These filters were implemented in the spatial domain and post-processed with a 31x31 variance estimator.

3.3 The Error Metric

The output of each filter/texture combination was plotted against illuminant tilt angle τ. According to the preceding theory each of these graphs should be a sinusoidal function of 2τ (equation 9). As we do not have the surface height functions we cannot directly calculate the parameters of these equations. We therefore compute the best fit sinusoid for each graph using Fourier analysis for the angular frequency 2 radians/radian [4] . To determine how close a dataset is to our theoretical prediction we calculate an error metric: the mean squared difference between the experimental data and their corresponding best-fit estimates.

[4] As the x-axis is the tilt angle which may be measured in either radians or degrees, the period may be expressed as either π radians or 180°. Corespondingly the angular frequency may be expressed as $frac\pi 90$ radiansdegree or as 2 radiansradian.

3.4 Results (1D)

Figure 4 provides a graphical summary of our results. It is a histogram of the error metrics obtained from the experiments. It shows an asymmetric distribution in which errors range from 0.01, signifying an extremely good fit, up to three outliers at around 0.29. Figure 5 shows the behaviour of two complex Gabor filters over the complete range of thirty textures. It also contains the median, upper, and lower quartile lines (at 0.036, 0.056 and 0.024 respectively) determined over all of the results.

In order to provide an insight into what these error values mean we have selected twelve sample graphs for display. Four results closest to the median error, four results closest to the lower quartile, and the four closest to the upper quartile, have been selected to represent typical, good, and poor results respectively. These are shown in figure 6, figure 7, and figure 8 respectively.

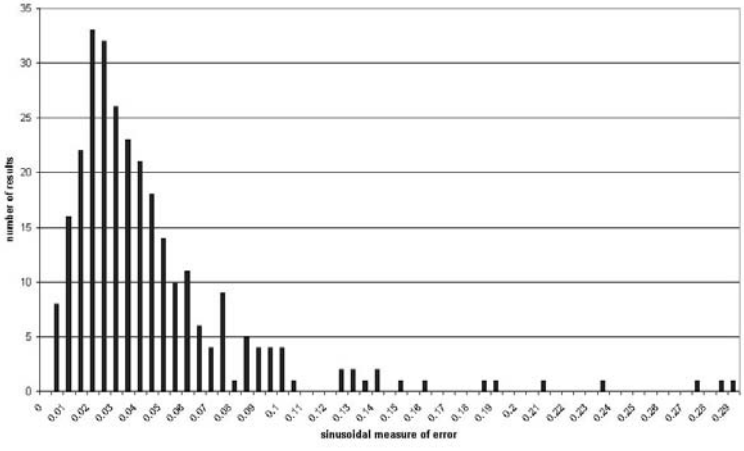

Fig. 4. Histogram of error metric values for the goodness of fit of sinusoidal functions to feature/tilt angle graphs. Low values, e.g. the eight occurrences of 0.01, indicate that very good sinusoidal fits were obtained for those graphs of feature output against illuminant tilt angle. High values, e.g. the three outliers around 0.28 indicate a bad fit

What is evident from these results is that even filter/texture combinations with 'poor' error metrics (shown in figure 8) follow the sinusoidal behaviour quite closely. This is all the more surprising considering how many of our textures significantly violate the 'no shadow' and 'low slope angle' assumptions. Textures *michael1* and *michael5* for instance obviously contain many shadows and high slope angles (see table 1) yet the tilt responses shown in figure 8 show that they are still approximately sinusoidal.

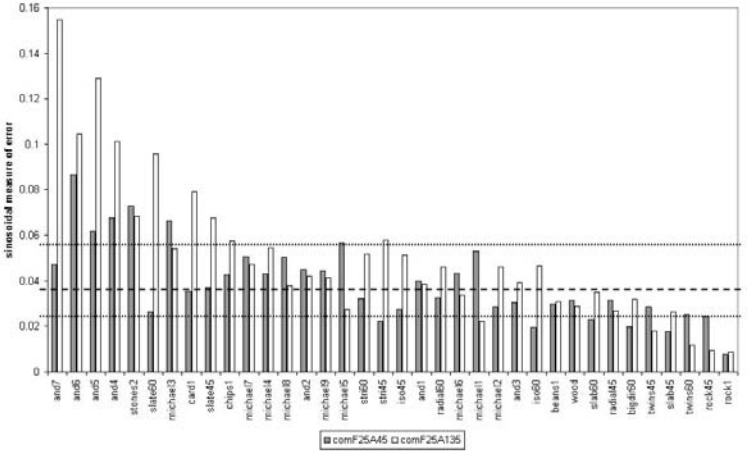

Fig. 5. Bar-chart of error metric values for the complex Gabor filters F25A45 & F25A135 showing how the goodness of fit varies over the different texture samples. The dashed lines show the median, upper and lower quartile values. Sample images of the textures *and7* to *rock1* are shown in the table 1.

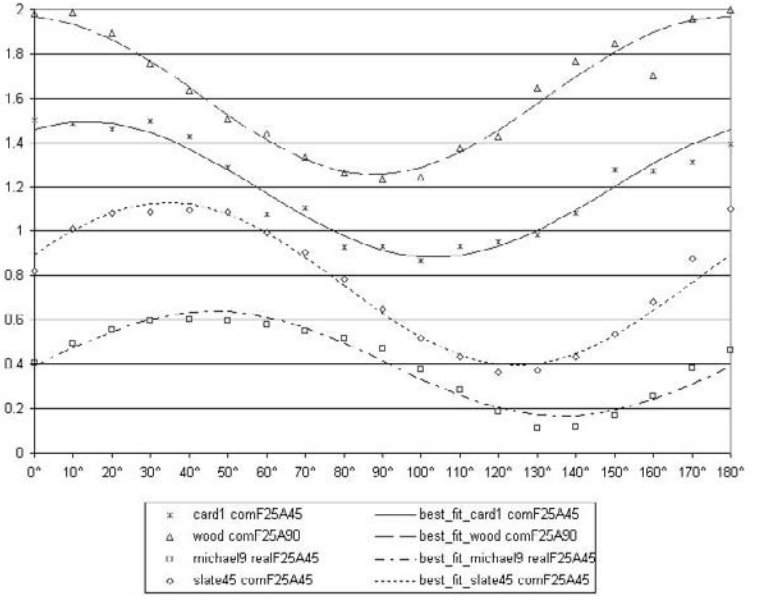

Fig. 6. Four datasets with error metrics closest to the **median error** of 0.036. (Each plot shows how one output of one feature varies when it is repeatedly applied to the same physical texture sample, but under varying illuminant tilt angles. Discrete points indicate measured output and the curves show the best-fit sinusoids of period 2τ)

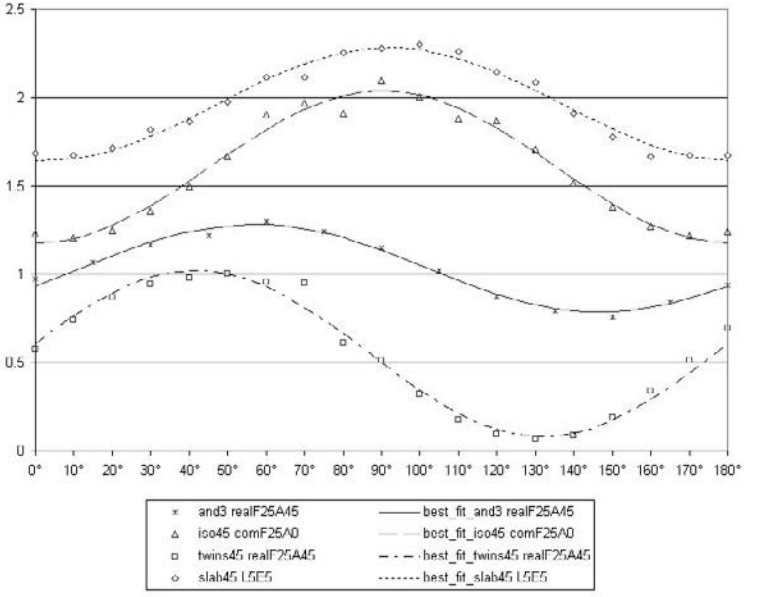

Fig. 7. The four datasets closest to the lower quartile (0.024)

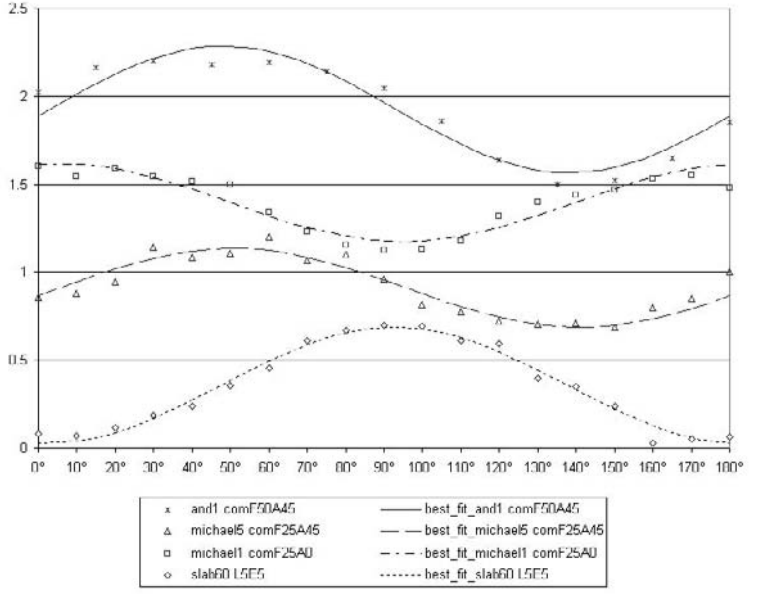

Fig. 8. The four datasets closest to the upper quartile (0.056)

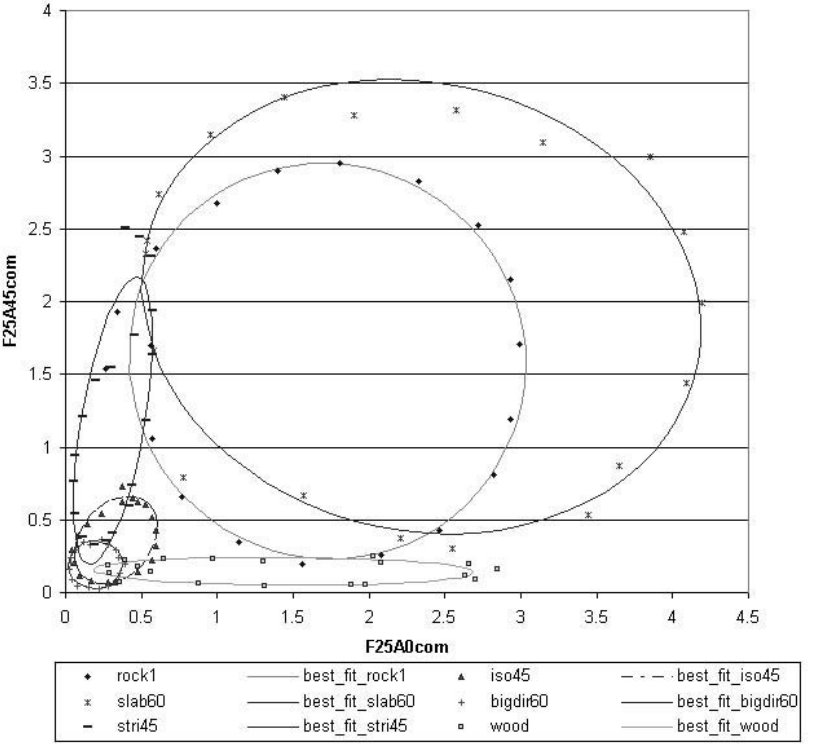

Fig. 9. The behaviour of six textures in the comF25A0/comF25A45 feature space together with the best fit ellipses (each point on an ellipse denotes a different value of illuminant tilt)

3.5 Results (2D)

Given that the 1D results show that a large number of feature/texture combinations follow a sinusoidal function of tilt, then the behaviour of cluster means in a multi-dimensional feature space should be of a super-elliptical form. In order to provide additional evidence for this proposition we plotted the results in a number of 2D feature spaces. We used the parameters derived from the best-fit sinusoids to determine the equations of the ellipses in these 2D spaces.

Figure 9 shows one of these sets of plots. The behaviours of cluster means of six textures in the F25A0/F25A45 feature space are shown together with their best fit ellipses. Because in this case the filters are identical except a difference in direction of 45°, the curve of the isotropic surface *rock1* is a circle.

It can be seen that the ellipses of different textures cross each other showing that illumination tilt invariant classification using a simple linear classifier is not possible.

Table 1. One image of each of the thirty sample textures

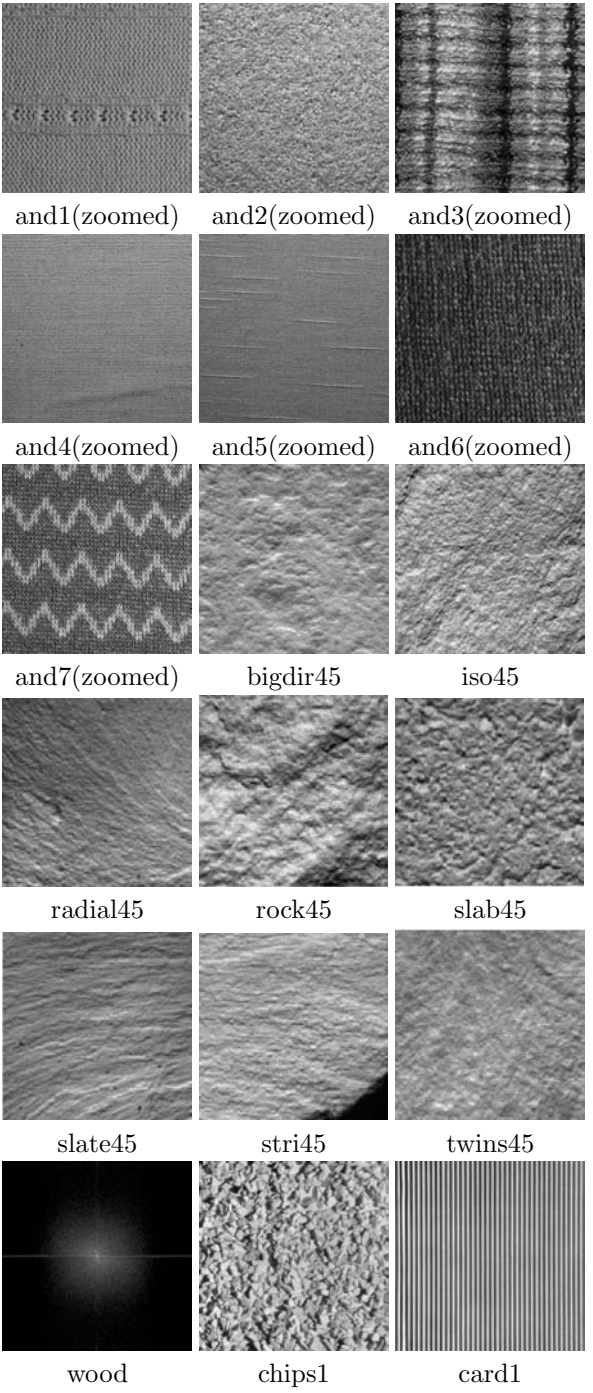

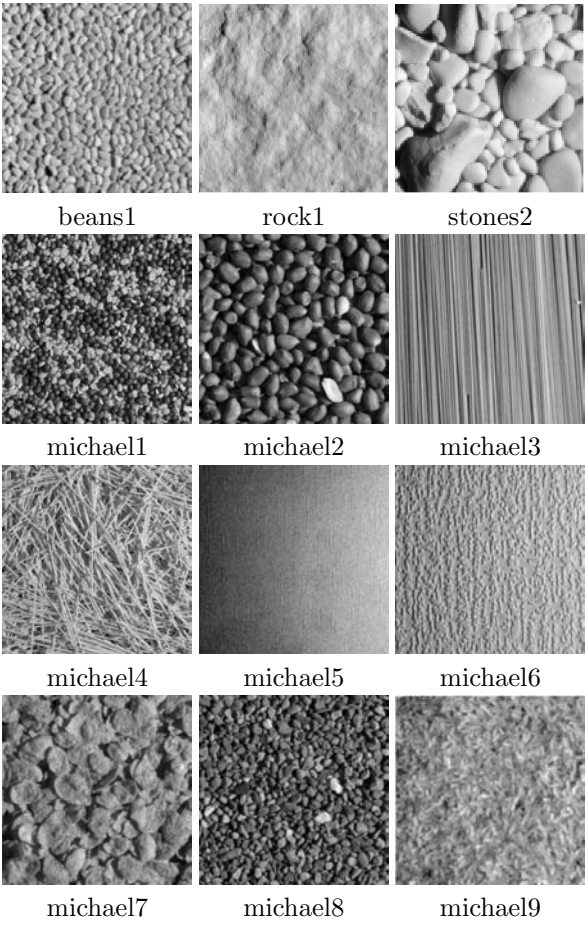

4 Conclusions

We have presented new theory and new experimental results that support our proposal that the means of texture features derived from linear filters are sinusoidal functions of illuminant tilt (τ). That is:

$$f(\tau) = a + d\cos(2\tau + \phi)$$

As these behaviours are of the same frequency, but differ in amplitude and phase, they form a set of simple harmonic components. The trajectories of cluster means therefore describe Lissajous's figures in a two dimensional feature space, and in general follow super-elliptical trajectories in multidimensional feature spaces.

What is perhaps most surprising is that the empirical results agree well with these theoretical predictions given that many of our samples clearly violate the low-slope-angle and no-shadow assumptions.

These results explain why texture classifiers can fail when the tilt angle is changed and suggest that, given suitable training data, this behaviour may be exploited for the design of tilt invariant texture classifiers.

References

1. M.J. Chantler. *Why Illuminant Direction is Fundamental to Texture Analysis*, IEE Proceedings on Vision, Image and Signal Processing, Vol. 142", No.4, August, 1995 pages 199-206
2. P. Brodatz. *Textures: a photographic album for artists and designers.* Dover, New York, 1966.
3. P.R. Kube and A.P. Pentland. On the imaging of fractal surfaces. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 10(5):704–707, September 1988.
4. T. Randen and J.H. Husoy. Filtering for texture classification: A comparative study. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(4):291–310, April 1999.
5. M.J. Chantler and G. McGunnigle. The Response of Texture Features to Illuminant Rotation *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* Vol III: pages 955–958, 2000.
6. K.J. Dana, S.K. Nayar, B. van Ginneken, and J.J. Koenderink. Reflectance and texture of real-world surfaces. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 151–157, 1997.
7. K.J. Dana and S.K. Nayar. Histogram model for 3d textures. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 618–624, 1998.
8. B. van Ginneken, J.J. Koenderink, and K.J. Dana. Texture histograms as a function of irradiation and viewing direction. *International Journal of Computer Vision*, 31(2/3):169–184, April 1999.
9. K.J. Dana and S.K. Nayar. Correlation model for 3d texture. In *Proceedings of ICCV99: IEEE International Conference on Computer Vision*, pages 1061–1067, 1999.
10. T. Leung and J. Malik. Recognizing surfaces using three-dimensional textons. In *Proceedings of ICCV99: IEEE International Conference on Computer Vision*, pages 1010–1017, 1999.
11. T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *International Journal of Computer Vision*, 43(1):29–44, June 2001.
12. A.K. Jain and F. Farrokhnia. Unsupervised texture segmentation using gabor filters. *Pattern Recognition*, 24(12):1167–1186, December 1991.
13. K.I. Laws. *Textured Image Segmentation.* PhD thesis, Electrical Engineering, University of Southern California, 1980.