Rotation Invariant Classification of 3D Surface Textures using Photometric Stereo and Surface Magnitude Spectra

M. J. Chantler & J. Wu
Department of Computing and Electrical Engineering
Heriot-Watt University
Edinburgh, EH14 4AS, Scotland
http://www.cee.hw.ac.uk/~mjc

Abstract

Many image-rotation invariant texture classification approaches have been presented. However, image rotation is not necessarily the same as *surface* rotation. This paper proposes a novel scheme that is surface-rotation invariant. It uses magnitude spectra of the partial derivatives of the surface obtained using photometric stereo. Unfortunately the partial derivative operator is directional. It is therefore not suited for direct use as a rotation invariant feature. We present a simple frequency domain method of removing the directional artefacts. Polarograms (polar functions of spectra) are extracted from resulting spectra. Classification is performed by comparing training and classification polarograms over a range of rotations (1° steps over the range 0° to 180°). Thus the system both classifies the test texture and estimates its orientation relative to the relevant training texture.

A proof for the removal of directional artefacts from partial derivative spectra is provided. Results obtained using the classification scheme on synthetic and real textures are presented.

1 Introduction

Many texture classification schemes have been presented that are invariant to image rotation [1,2,3]. They normally derive their features directly from a single image and are tested using rotated images. If the image texture results solely from albedo variation rather that surface relief, or if the illumination is not directional or immediately overhead; then these schemes are surface-rotation invariant as well. However, in many cases *rotation of a textured surface produces images that differ radically* from those provided by pure image rotation. This is mainly due to the directional filtering effect of imaging using side-lighting [4, 5].

We present a highly novel *surface rotation invariant* approach to texture classification. Our approach uses polarograms [6] derived from surface derivative spectra. We use photometric stereo to obtain the required partial derivative fields. They are Fourier transformed and combined to provide a frequency domain function that does not contain the directional artefacts associated with partial derivatives. Polarograms of this function are compared with those of training classes using a goodness-of-fit measure to provide rotation invariant texture classification.

1.1 Organisation of this paper

The next section describes the problem and related work. Section 3 presents theory that shows that our polarograms do not suffer from directional artefacts. Section 4 describes the complete classification scheme, while section 5 presents results obtained by applying this scheme to synthetic and real textures.

2 The problem

This section discusses the problems that occur when 3D surface textures are imaged using directional illumination. First we define three terms:

3D surface texture - the three-dimensional variation of surface relief,

Albedo texture - the of variation of the albedo property due, for instance, to a printed 'pattern' on a flat surface, and

Image texture - the variation in grey-scale or colour value in an image that results from imaging 3D surface texture, albedo texture or a combination of the two.

3D surface textures are often illuminated from one side when they are photographed in order to enhance their image texture, e.g. [7]. Such imaging acts as a *directional filter* of the 3D surface texture [5]. This phenomena is illustrated below. Figure 1 shows the same sample of 3D surface texture imaged using two different illumination conditions.

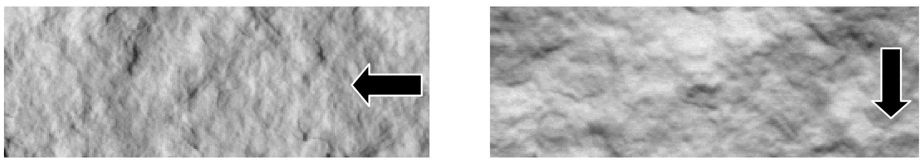

$\tau = 0°$ $\tau = 90°$

Figure 1 - Images of the same 3D surface texture imaged at two different illuminat tilt (τ) directions (shown by the arrows)

The directional filtering effect is more obvious in the frequency domain. Figure 2 shows the Fourier transforms of the two image textures shown above.

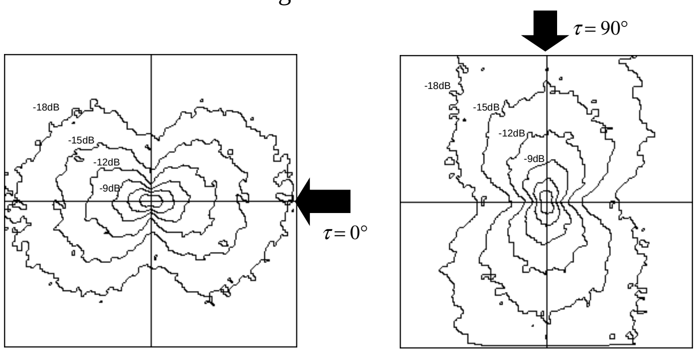

Figure 2 - FFTs of the images shown in Figure 1

This shows more clearly that side-lighting accentuates texture components in the illuminant tilt direction (τ) and attenuates those at right angles. It follows that if the surface is rotated and imaged under directional illumination conditions, the images produced will be a function of the rotated surface's characteristics *and the directional filtering effect* described above. The four image textures below illustrate this point.

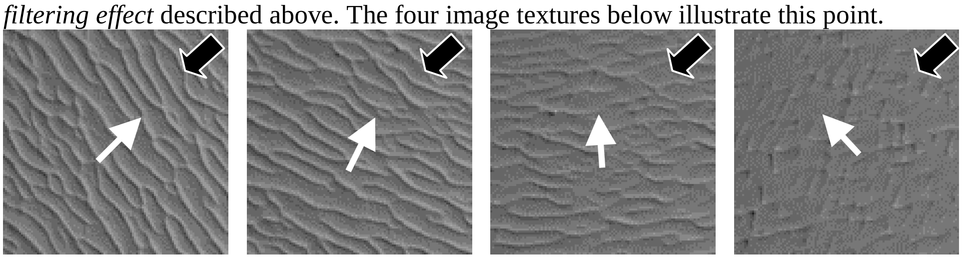

$\varphi = 45°$ $\varphi = 75°$ $\varphi = 105°$ $\varphi = 135°$

Figure 3 - four images of a directional texture. The *surface* has been rotated at angles of 45°, 75°, 105°, and 135° (white arrows). The illuminant tilt was kept constant at $\tau = 45°$ (black block arrows). Note that when the 3D surface texture direction is at right angles to the illuminant tilt the texture is hardly discernable.

These images show that rotation of a 3D surface texture does not result in a simple rotation of the image texture. Conventional image rotation invariant algorithms [1,2,3] do not take this important and everyday phenomena into account.

2.1 Related work

As previously stated many texture classification schemes have been presented that are invariant to image rotation [1,2,3]. Few take into account the problems caused by illumination described above. Exceptions include Leung and Malik's classification system which is trained on textures that are each imaged under 20 different illumination and orientation conditions [8]. This generalises the classifier but does not use explicit 3D surface texture information directly. Their textures were obtained from the Columbia-Utrecht Reflectance and Texture Database [9]. Dana and Nayer describe a correlation model for 3D surface texture and suggest how this might be used to provide a 3D surface texture feature, correlation length. They do not however, use this for texture classification purposes [10].

McGunnigle and Chantler proposed scheme that used photometric stereo to obtain gradient information, these data were used to synthesise images of the training textures once the classification illumination conditions were known [11, 12]. The scheme was not rotation invariant. Later they proposed another photometric-based system, however, this time the gradient information was directly filtered using isotropic gabor filters to provide a rotation insensitive scheme [13].

Smith also uses 3D surface texture information directly [14]. He uses photometric stereo to acquire surface gradient information and suggests the use of features derived from the gradient space (including attitude, principal orientation, shape factor, and shape distribution) for the 'quantitative analysis of repetitive surface textures'. He does not go as far as applying this approach to the task of classification of rough surfaces using a conventional classifier - although it would be very interesting to see the results. In many respects our own gradient space approach which uses eigenvalues as the features is very similar [15]. However, this paper presents a different approach; it uses frequency domain information rather than gradient space features.

3 Theory

Our approach uses estimated partial derivatives of the surface height function. We obtain these using photometric stereo [12, 16, 17]. Unfortunately the partial derivative operator is a directional filter of surface texture. This means that the partial derivatives cannot be used directly in a rotation invariant classifier; they must be processed first.

This section proposes a method by which the partial derivatives may be combined in the frequency domain in such a way as to remove these directional artefacts. The gradient estimates provided by photometric stereo are normally in the form of the partial derivative fields $p(x,y)$ and $q(x,y)$.

$$p(x,y) = \partial z / \partial x \tag{1}$$

$$q(x,y) = \partial z / \partial y \tag{2}$$

where: $z(x,y)$ is the surface height function of a texture in the x-y plane.

The Fourier transforms of (1 & 2) are:

$$P(u,v) = iuS(u,v) = i\omega(\cos\theta)S(\omega,\theta) \tag{3}$$

$$Q(u,v) = ivS(u,v) = i\omega(\sin\theta)S(\omega,\theta) \tag{4}$$

where: $S(u,v)$ and $S(\omega,\theta)$ are the Fourier transforms of surface height function in cartesian and polar forms,
u,v are spatial frequency variables, and
ω,θ are their polar equivalents.

Now (3 & 4) show that both derivatives act as directional filters due to the $\cos\theta$ and $\sin\theta$ terms. However, we may combine them to provide a function free of directional artefacts:

$$M(\omega,\theta) = |P(u,v)|^2 + |Q(u,v)|^2 = [\ \omega|S(u,v)|\ ]^2 \tag{5}$$

To utilise the directional information contained in $M(u,v)$ we use its polarogram:

$$\Pi(\theta) = \int_0^{\infty} M(\omega,\theta)\,d\omega \tag{6}$$

Note that while both $M(\omega,\theta)$ and its polarogram do not theoretically contain any directional artefacts they are however, rotationally sensitive. That is, if the surface is rotated by an angle φ then a new spectrum $M_\varphi(\omega,\theta)$ and a new polarogram $\Pi_\varphi(\theta)$ will result (see Figures 6 & 9), such that:

$$\Pi_\varphi(\theta) = \Pi(\theta + \varphi) \tag{7}$$

Hence in order to provide rotationally invariant discriminate functions we compare the test texture's polarogram against each of the training textures' polarograms. This comparison is performed over a range angular displacements of the test polarogram (φ_{test} = 0°, 1°, 2°, ….180°).

4 Surface Rotation Invariant Classification

The complete classification scheme is shown in Figure 4. A photometric image set (three images taken at illuminant tilt angles of 0°, 90° and 180° of the test texture) is captured. The photometric algorithm uses these to provide estimates of the partial derivative fields p and q. These are Fourier transformed and processed to provide an estimate of the texture's polarogram, which is compared with polarograms obtained from training images over a range of angular displacements (φ_{test}). In fact in order to exploit some of the radial information we actually take four polarograms of each texture over different frequency bands (ω = 1-4, 5-8, 9-12, and 13-16 cycles per image width). We use a sum of squared differences as a goodness-of-fit indicator at each displacement value (φ_{test}) and choose the best fitting combination of displacement (φ_{test}) and training texture to provide both classification output and an orientation estimate for the test texture.

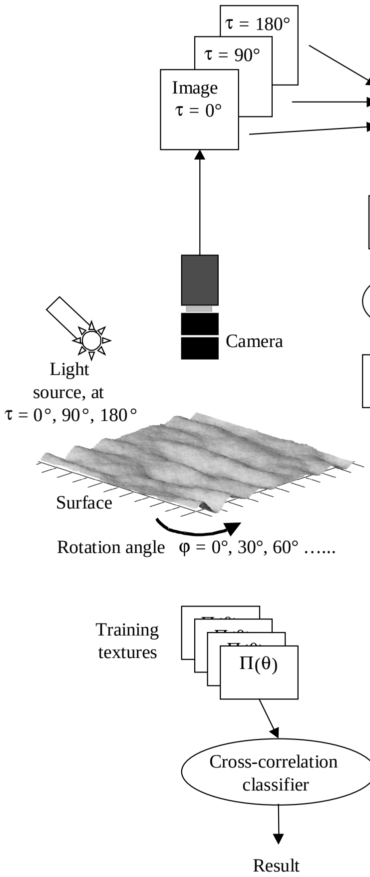

Figure 4 - The surface rotation invariant classification scheme

5 Experiments

A set of experiments were performed on synthetic and real textures. The experiments on the synthetic textures were designed to show that the $M(\omega,\theta)$ functions and their polarograms are rotationally sensitive but contain no directional artefacts. That is $M_{\varphi}(\theta) = M(\theta + \varphi)$ and similarly $\Pi_{\varphi}(\theta) = \Pi(\theta + \varphi)$. These textures were also used to test the basic classifier. Further testing was performed using four real textures.

5.1 Synthetic Textures

Montages of a selection of synthetic images (*Rock, Sand, Ogil, Malv*) are shown in Figure 5. These textures are defined in [12]. From Figure 6 it can be seen that the $M(\omega,\theta)$ functions of the rotated textures are simply a rotation of the original (φ = 0°) $M(\omega,\theta)$ function as predicted. This shows that the $M(\omega,\theta)$ function conatins no directional artefacts such as a directional filtering effect.

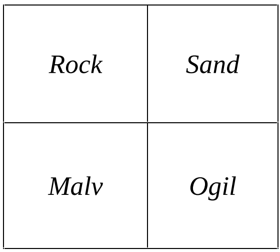

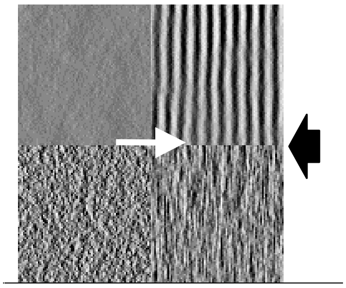

Rotation angle φ = 0°

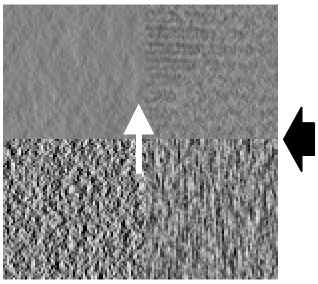

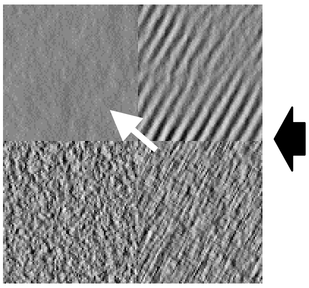

Rotation angle φ = 90° Rotation angle φ = 150°

Figure 5 - Synthetic textures at three surface rotations with constant illumination. Surfaces, rotated as indicated by white arrows, rendered at an illuminant tilt of 90° (black arrows) and combined into montage for display purposes.

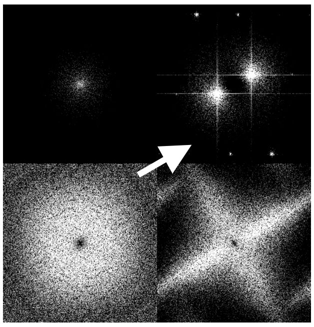

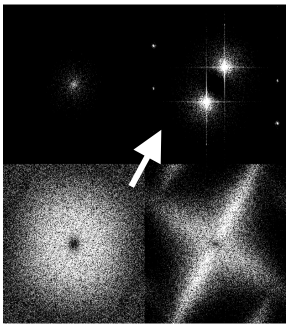

Rotation angle φ = 30° Rotation angle φ = 60°

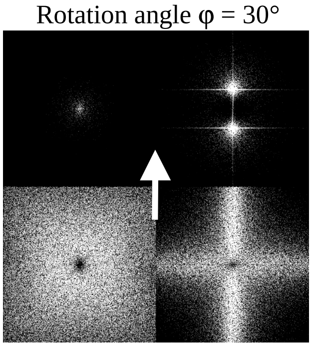

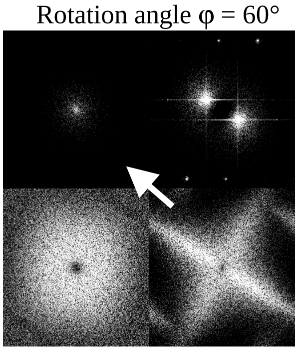

Rotation angle φ = 90° Rotation angle φ = 150°

Figure 6 – $M(\omega,\theta)$ functions of <u>each</u> of the 4 synthetic textures shown in montage format (as above) for 4 surface rotations.

It is not surprising therefore, that Figure 7, dipicting the polarograms of the rotated *Sand* surface, shows that a rotated texture's polarogram is an approximate translation of the unrotated texture's polarogram.

For classification the system was trained using photometric image sets synthesised from the four surfaces obtained at a surface rotation angle of φ = 0°. This provided the four *'training'* polarograms (one for each class). The *classification* was performed on 256x256 samples of the textures rotated at 30° increments over the range 0° to 180°. Each of the resulting 76 polarograms was compared with each of the 4 'training' polarograms. The class was assigned to training texture with the best-fit polarogram. No classification failures occurred. A more testing experiment is described below.

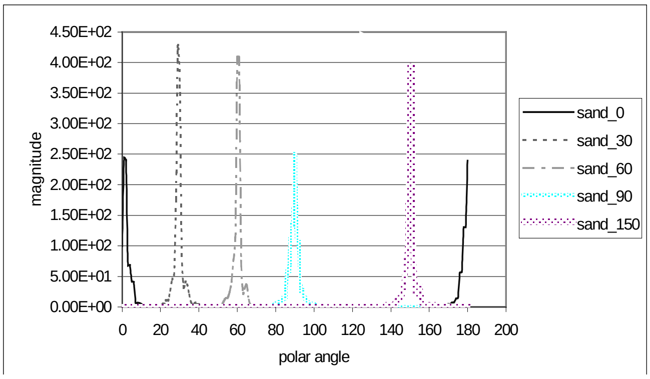

Figure 7 - Polarograms of the synthetic *Sand* texture at surface rotations of φ = 0°, 30°, 60°, 90°, and 150°.

5.2 Real Textures

Photometric image sets of four real textures were obtained over a range of surface rotations (0° to 180° in 30° steps). Montages of a selection of these images are shown below (Figure 8).

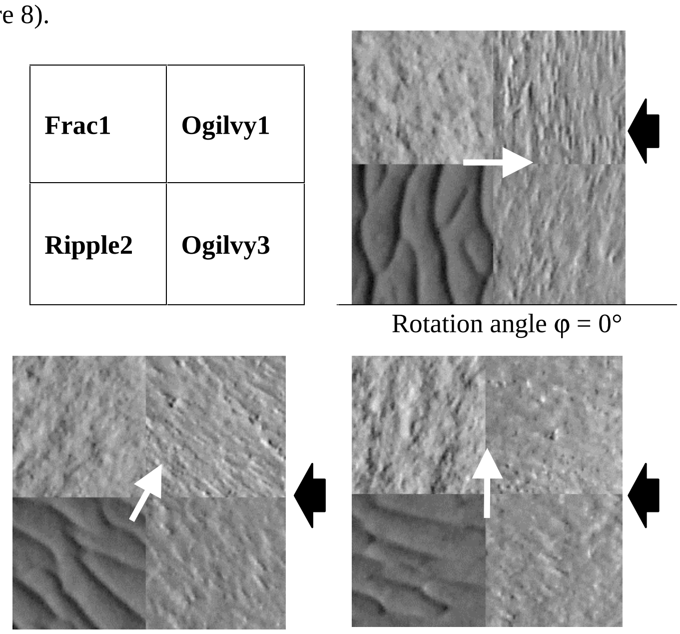

Rotation angle φ = 0°

Rotation angle φ = 60°

Rotation angle φ = 90°

Figure 8 – Real textures shown in montage format

For illustration purposes montages have been constructed from single images of the textures. The textures were in fact rotated individually. All textures shown above were captured at an illuminant tilt angle $\tau = 0°$. Images were also captured using iluminant tilt angles 90° and 180° images as well but are not shown here).

Comparison of the images *Ogilvy1*(φ = 0°) and *Ogilvy1*(φ = 90°), show that they are not simple rotations of each other. In *Ogilvy1* (φ = 0°) the vertical lines of the texture are clearly present. In *Ogilvy1*(φ = 90°) the image texture has been attenuated and the lines (which should be vertical) are no longer visible. Similarly for *Frac1*'s images.

Figure 6 below shows the polarograms of the *Ogilvy1* texture.

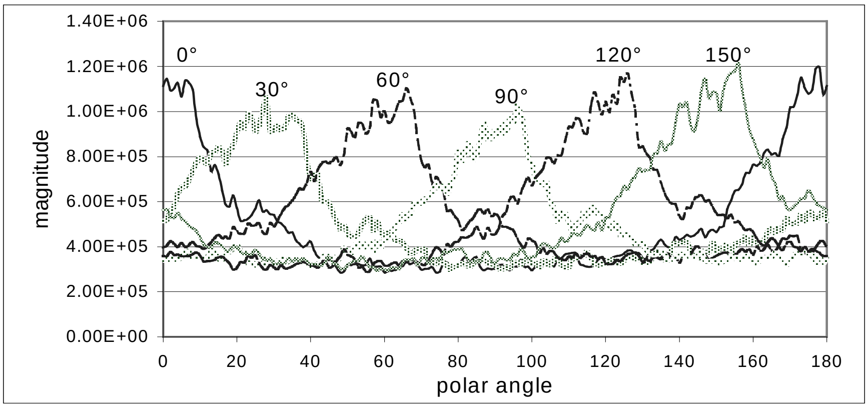

Figure 9 - Polarograms of the real *Ogilvy1* texture at surface rotations of $\varphi = 0°$, 30°, 60°, 90°, 120° and 150°.

It can be seen that these polarograms are again approximate translations of each other, and that the degree of each translation approximates the corresponding rotation of the surface.

Figure 7 shows polarograms of all four textures obtained at a surface rotation angle of $\varphi = 0°$, these form the 'training' polarograms.

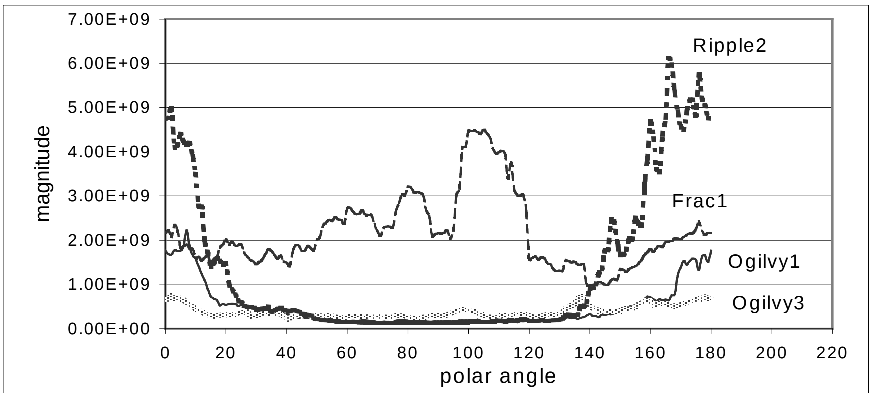

Figure 10 - Training polarograms (for a 256x256 sample size)

The classification task therefore becomes one of comparing the polarogram of the test texture against each of the polarograms shown in Figure 10. Again the test polarogram must be 'rotated' (or rather translated) to find the best-fit. The texture is then assigned to the class for which the lowest sum of square differences occurs at each of these best-fit rotation angles. This classification test was repeated using sample sizes 128x128, 64x64 and 32×32. The samples were cut from the original 256x256 images to provide the maximum number of samples without overlap. Hence for the 32x32 case 4864 classifications were performed (64 samples per rotation x19 rotations x 4 textures). The results are shown below in Table 1 and Figure 11.

	32*32	**64*64**	**128*128**	**256*256**
Frac1	46.9%	74.1%	78.6%	100%
Ogilvy1	34.4%	55.4%	67.9%	100%
Ogilvy3	87.1%	86.6%	96.4%	100%
Ripple2	59.2%	93.8%	96.4%	100%
No. of classifications	4864	1216	304	76
Average of accuracy	57%	77%	85%	100%
Misclassification rate	43%	23%	15%	0%

Table 1 - Classification results for real textures

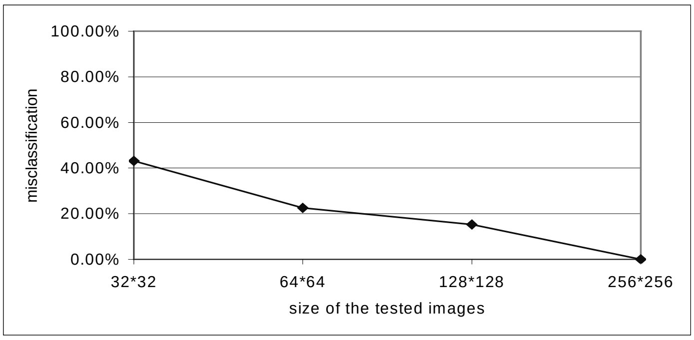

Figure 11 – Misclassification rates for real textures.

The results show that the approach is very sensitive to the sample size - this is probably due to the errors that result from computing the polarogram (a polar measure) from small discrete cartesian FFTs. However, we have not had time to investigate this further.

Nevertheless the results show that for the larger sample sizes the approach attains remarkable accuracy for a set *real and rotated 3D surface textures*.

6 Conclusions

1. This paper has proposed a completely novel texture classification scheme that is surface-rotation invariant.
2. We have presented theory and experiment that shows that the basic feature set, polarograms derived from surface gradient spectra, is free of directional artefacts. We have shown that this is *not* the case for single side-lighted images of surface texture.

3. We have also presented results that show that our classification scheme achieves very good classification rates using large sample sizes and poor classification rates using smaller sizes.
4. We therefore think that this classification scheme is suited to classification of sample images and the estimation of the texture direction, but that a gradient space approach such as that presented in [15] would be better for use on texture segmentation problems.

References

[1] R. Porter & N. Canagarajah, "Robust rotation invariant texture classification: Wavelet, Gabor filter and GMRF based schemes", IEE Proc. Vis. Image Signal Process. Vol.144, No.3, June 1997.

[2] F.S. Cohen, Z. Fan & M.A.S. Patel, "Classification of rotated and scaled textured images using Gaussian Markov field models", PAMI V13, February 1991, pp192-202

[3] J. Mao & A.K. Jain, "Texture classification and segmentation using multiresolution simultaneous autoregressive models", Pattern recognition, V25, No.2, 1992, pp 173-188

[4] M.J. Chantler, "The effect of illuminant direction on texture classification", PhD Thesis, Dept. of Computing and Electrical Engineering, Heriot-Watt University, 1994.

[5] M.J. Chantler, G.T. Russell, L.M. Linnett "Illumination: A directional filter of texture?", British Machine Vision Conference 1994 Vol.2 pp.449-458.

[6] L.S. Davis, "Polarograms : a new tool for image texture analysis", Pattern Recognition, V13, No. 3, 1981, pp219-223.

[7] P. Brodatz, "Textures: a photographic album for artists and designers", Dover, New York, 1966.

[8] T. Leung and J. Malik, "Recognizing Surfaces using Three-Dimensional Textons", IEEE International Conference on Computer Vision, Corfu, Greece, Sep 1999

[9] K.J. Dana, B. van Ginneken, S.K. Nayar and J.J. Koenderink, "Reflectance and Texture of Real World Surfaces", ACM Transactions on Graphics, Volume 18, No. 1, pp. 1-34, January 1999.

[10] K.J. Dana, S.K. Nayar, " Correlation Model for 3D Textures", IEEE International Conference on Computer Vision, Corfu, Greece, Sep 1999

[11] G. McGunningle, M.J. Chantler, "A model-based technique for the classification of textured surfaces with illuminant direction invariance", Proceedings of BMVC97, Essex, October 1997, pp470- 479.

[12] G. McGunnigle, "The classification of textured surfaces under varying illuminant direction", PhD Thesis, Dept. of Computing and Electrical Engineering, Heriot-Watt University, 1998.

[13] G. McGunnigle, M.J. Chantler, "Rotation insensitive classification of rough surfaces", IEE Proceedings: Vision, Image and Signal Processing, Vol. 146, No. 6, December 1999, pp345-352.

[14] M. L. Smith, "The analysis of surface texture using photometric stereo acquisition and gradient space domain mapping", Image and Vision Computing Journal, V17, 1009-1019, 1999.

[15] G. McGunnigle M.J. Chantler, "Rough surface classification using point statistics from photometric stereo", to appear in Pattern recognition Letters.

[16] G. Kay and T. Caelli, "Estimating the Parameters of an Illumination Model Using Photometric Stereo", Graphical Models and Image Processing, Vol.57, No.5 Sept. 1995, pp.365-388.

[17] R. Woodham, "Photometric method for determining surface orientation from multiple images", Optical Engineering, Jan./Feb. 1980, Vol.19 No.1, pp.139-144.