Rotation invariant classification of rough surfaces

G.McGunnigle and M.J.Chantler

Abstract: Rotation of a rough, textured surface will not produce a simple rotation of the image texture. It follows that where image texture is a function of surface topography, existing rotation invariant texture classification algorithms are not robust to surface rotation. The effect of surface rotation on the observed image is analysed using an existing theory, a novel scheme to stabilise classification accuracy is proposed and evaluated. The scheme uses photometric stereo to estimate the surface derivatives, which are then used as the input to a classifier. Simulations indicate that, where the level of image noise is moderate or low, the approach is successful in maintaining classification accuracy. Furthermore, in some circumstances, the extra information used by the algorithm allows classification accuracy superior to that based on one image alone, even without rotation.

1 Introduction

The development of texture classification algorithms that are robust to texture rotation is an important area of research within texture analysis, and dates back to the earliest days of the field [1]. Most rotation invariant schemes are derived from rotation sensitive algorithms, such as co-occurrence methods [1, 2], autoregressive techniques [3, 4], Markov random fields [5, 6], wavelets [6, 7], and Gabor filters [6, 8]. A brief review of rotation invariant algorithms is given in [6].

The algorithms can be divided into two classes: rotation insensitive and rotation invariant algorithms [9]. The former ignores directional information completely, whereas the latter exploits information about the orientation of texture components relative to each other while ignoring their absolute orientation. Rotation insensitive schemes include the circular autoregressive technique [3] and the isotropic Gabor filter used by Porter and Canagarajah [6]. Rotation invariant techniques are more complex and include techniques such as Kashyap's additional feature and Choe's two-level system [10]. This paper is restricted to rotation insensitive classifiers. The problem that is identified, however, occurs prior to the application of any algorithm and so is equally relevant to rotation invariant classifiers.

All of the algorithms discussed above deal with the rotation of image texture. However, in this paper it is argued that this is an insufficient model for many applications. Rotation of a rough surface is not equivalent to the rotation of its image. It follows that rotation invariant algorithms will not be robust to surface rotation.

This paper uses simulation to demonstrate the visual effect of rotating a rough surface, and analyses the effect using theory. The inability of an image rotation insensitive system to deal with surface rotation is also demonstrated. Hence the main objectives of this paper are: first, to explicitly show that rotation of a rough surface is not equivalent to rotation of its image, and that conventional rotation-invariant algorithms may fail under these conditions; and secondly, to propose a new classification scheme that uses surface gradient information derived from photometric stereo to provide a surface rotation invariant classifier.

IEE Proceedings online no. 19990707

DOI: 10.1049/ip-vis:19990707

Paper first received 1st December 1998 and in revised form 14th July 1999

The authors are with the Department of Computing and Electrical Engineering, Heriot-Watt University, Riccarton, Edinburgh EH14 4AS, UK

2 Problem description

2.1 Image effects

Consider the hypothetical inspection task shown in Fig. 1. A rough surface comprising two or more classes of surface topography is illuminated and is imaged using a monochrome CCD camera. The test piece may be rotated by an arbitrary angle ϕ from the angle at which training occurred. The position of the camera and illumination will be held constant. We shall assume the surface's reflectance characteristics to be uniform and diffuse throughout the surface, and the observed texture to be a product of the interaction of the illuminant and the surface topography. We will also assume an orthographic projection, that is, the image corresponds to a uniform sampling of intensity over the surface plane.

If the surface texture is isotropic, surface rotation will have no effect on the characteristics of the imaged texture (Fig. 2), though it is possible to follow the rotation of individual surface structures from frame to frame. If, on the other hand, the surface is directional, surface rotation will alter the imaged texture beyond simple rotation (Fig. 3). In this case, as the surface is rotated the directionality of the

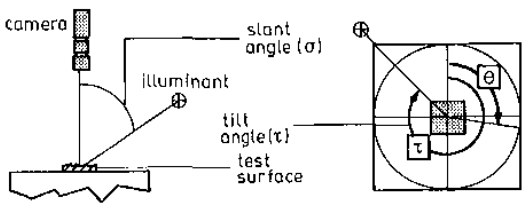

Fig. 1 *Geometry of a hypothetical inspection system*

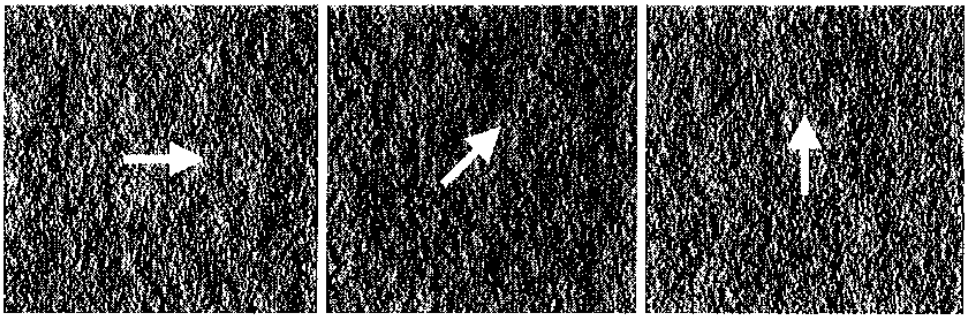

Fig. 2 *Images of isotropic surface rotated by 0°, 45° and 90°*
Arrow shows relative orientation of the surface

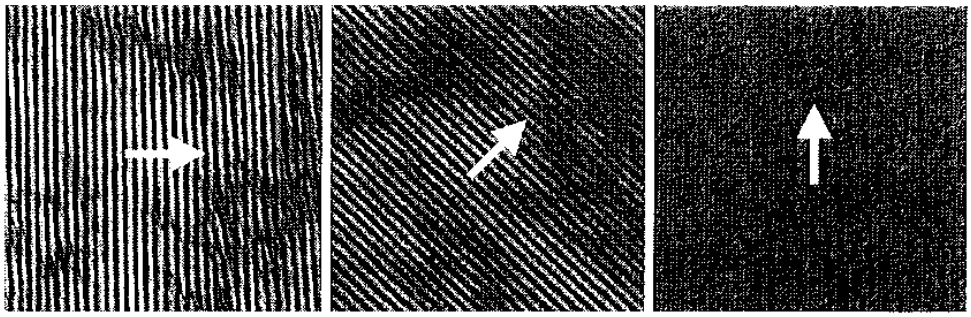

Fig. 3 *Images of directional surface rotated by 0°, 45° and 90°*
Arrow shows relative orientation of the surface

texture rotates. However, as the surface directionality (originally collinear with the illuminant) is rotated relative to the illuminant, the directionality visible in the image is attenuated. This gives the rotated surface an appearance distinct from that of the unrotated surface.

2.2 Classification

We begin with a brief description of an image rotation insensitive (IRIS) classifier. The feature set used in this paper consists of a set of isotropic Gabor filters, similar to those described in [6], though in this case the classifier will be used to segment, rather than classify the image. The filters measure the signal energy held within a radial frequency band; no directional filtering takes place and so the features contain no information about the directionality of the texture. In this way the features are unaffected by rotation of the data set.

The spectral response of a filter centred at ω_0 is shown in eqn. 1

$$G(\omega) = \exp\left(-\frac{(\omega - \omega_0)^2}{\sigma_\omega^2}\right) \tag{1}$$

where $G(\omega)$ is the frequency response of the filter, ω is the radial frequency, ω_0 is the centre frequency of the filter, and σ_ω is the bandwidth parameter of the filter.

Classification is carried out on the basis of the responses of four filters with centre frequencies at 64, 32, 16 and 8 cycles per image. The bandwidths of the filters, σ_ω, are 32, 16, 8 and 4 cycles per image, respectively. The power spectra of the filters used in this paper are shown in Fig. 4. The filtered images are then rectified and lowpass filtered to produce a set of feature images. These features are passed to a statistical classifier [11, p.125] and the image is segmented into regions using classification on a pixel-by-pixel basis. The process is summarised in Fig. 5.

The first objective of this paper is to show that image rotation-insensitive (IRIS) classifiers are not appropriate for surface rotation. To do this we must show that the classifier is: first, able to deal with image rotation; and secondly, unable to deal with surface rotation.

We define the test montage shown in Fig. 6, which consists of four textures which have been generated using a Lambertian rendering of synthetic height maps. The surfaces are assumed to be illuminated from a tilt angle of 270° and a slant angle of 50°. Clockwise from upper left-hand corner: surface 1 is an isotropic fractal surface with roll-off $\beta = 3.0$ and rms slope 0.25, surfaces 2 and 4 both exhibit fractal roll-off from a central spectral peak of radial frequency at 16 and 32 cycles per image, respectively. Surface 3 conforms to the model of directional surfaces proposed by Ogilvy [12] with vertical and horizontal correlation distances of 10 and 25 pixels,

Fig. 4 *Spectral bands of filters*

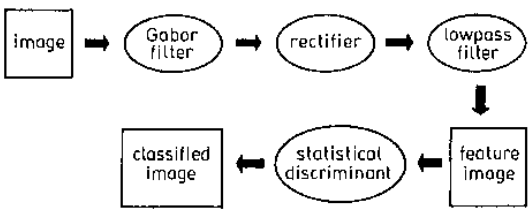

Fig. 5 *Classifier structure*

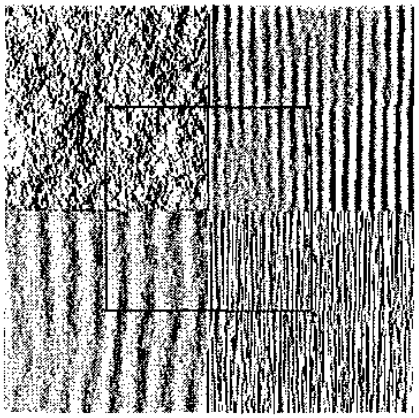

Fig. 6 *Test montage and training region*

respectively. The spectra of these surfaces are shown in Fig. 7.

The classifier is used for two tasks: in the first, the control experiment, the image is rotated. This is the test applied to conventional rotation invariant algorithms. In the second experiment the surface is rotated and is rendered under the same illumination geometry. In both cases the classifier is trained from the central area (25%) of the montage prior to rotation.

In the control experiment, the classifier is able to deal with image rotation with only a small fluctuation in classification accuracy (Fig. 8). The surface rotation experiment, in contrast, shows that the rotation invariant algorithm is unable to deal with the combined effect of the directional illuminant and surface rotation for a wide range of rotations. We conclude that image rotation is not equivalent to surface rotation and that image rotation invariant algorithms may not be suitable for the rotation invariant texture classification of rough surfaces.

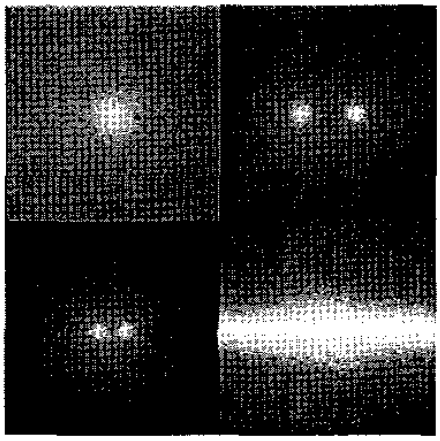

Fig. 7 *Log power spectra of test surfaces*

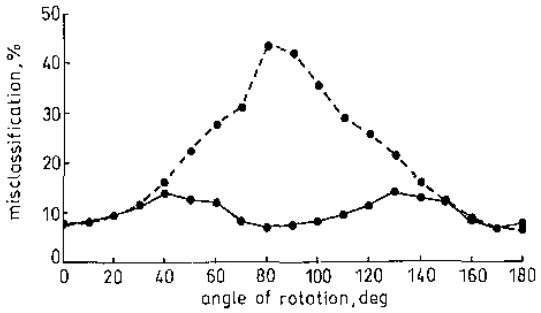

Fig. 8 *Accuracy of the IRIS classifier for image and surface rotation*

• · image rotation
– – • – surface rotation

3 Theory

3.1 Image

Kube and Pentland [13] developed a model, defined in spectral terms, which links the surface (which we wish to classify) to the image (on which we wish to make the classification). Kube and Pentland assume the surface reflectance to be Lambertian. The image $i(x, y)$ can then be expressed as a function of the illuminant orientation (τ, σ) and the surface partial derivative fields, $p(x, y)$ and $q(x, y)$.

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

A linear approximation (eqn. 3) to this equation is used to derive a frequency domain expression (eqn. 4). The assumption of linearity is valid for surfaces with moderate ($< 15°$) slope angles [13].

$$i(x,y) = -p(x,y)\cos\tau\sin\sigma - q(x,y)\sin\tau\sin\sigma + \cos\sigma \tag{3}$$

Since differentiation is a linear operation this can be transformed into the frequency domain, and expressed as a function of the surface magnitude spectrum

$$I_m(\omega,\theta) = -i\omega\sin\sigma[\cos(\theta)\cos(\tau) + \sin(\theta)\sin(\tau)]S_m(\omega,\theta-\phi) \tag{4}$$

where $I_m(\omega, \theta)$ is the image magnitude spectrum, and $S_m(\omega, \theta - \phi)$ is the magnitude spectrum of the surface rotated by an angle ϕ.

For the purposes of this paper it is more convenient to express this in terms of the power spectrum in the form shown in eqn. 5.

$$I(\omega,\theta) = \omega^2|\sin\sigma|^2|\cos(\theta-\tau)|^2 S(\omega,\theta-\phi) \tag{5}$$

where $I(\omega,\theta)$ is the image power spectrum, and $S(\omega, \theta - \phi)$ is the surface power spectrum rotated by an angle ϕ.

The imaging process may be thought of as a directional highpass filtering of the surface [14]. If we consider only the directional aspects of this filtering we can see that the image directionality is a product of the illuminant tilt angle (τ) as well as the surface directionality. Thus a surface rotation will not, in general, be equivalent to image rotation if the illuminant is not also rotated.

It is useful to focus on the directional aspects of the magnitude spectrum. We integrate the spectrum over radial frequency, over the range $\omega = 0.125\omega_s$ to $\omega = 0.75\omega_s$, where ω_s is the sampling frequency, to obtain a one-dimensional distribution of frequency components as a function of θ.

For an isotropic surface, image directionality is due solely to the directional effect of the illuminant. Rotation of the surface will cause no significant change in the image directionality so long as the illuminant direction is held constant (Fig. 9).

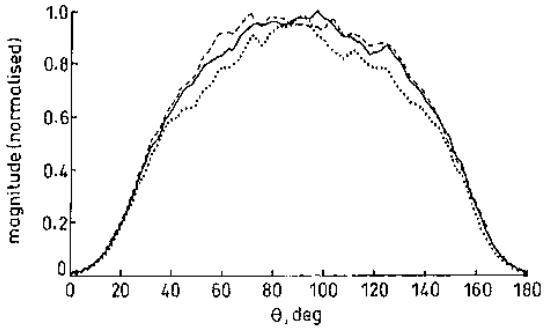

Fig. 9 *Polar plot of image spectrum from isotropic surface*
—— rotated by $\phi = 0°$
····· rotated by $\phi = 45°$
--- rotated by $\phi = 90°$

The directionality in the image of a directional surface is the product of both the surface and illuminant directionalities. The surface directionality will either be accentuated or attenuated depending on its orientation relative to that of the illuminant. Thus rotation of the directional surface causes the peak in the polar plot to be shifted. However, the amplitude of the peak falls as the surface directionality moves away from the angle of illumination, until, for $\phi = 90°$ the peak is completely attenuated (Fig. 10).

3.2 Features

The classifier uses features which are derived from the image. It follows that if the image is affected then the features, and ultimately the classifier, will also be affected. Although the nonlinear nature of classification limits our ability to analyse it, we can observe the effect of surface rotation on features, and so gain insight into the effect at the decision level of the classifier.

The isotropic Gabor filters act as isotropic bandpass filters. The power spectrum of the filtered images is given by

$$F(\omega, \theta) = |G(\omega)|^2 I(\omega, \theta) \tag{6}$$

Using eqn. 5 we may write

$$F(\omega, \theta) = |G(\omega)|^2 \omega^2 |\sin\sigma|^2 |\cos(\theta - \tau)|^2 S(\omega, \theta - \phi) \tag{7}$$

Although the filters are real, and therefore not admissible, we assume the filtered images have a mean value of zero. We achieve this by removing the mean component of the images prior to filtering. This of course, assumes that all the texture classes share the same mean. We justify this

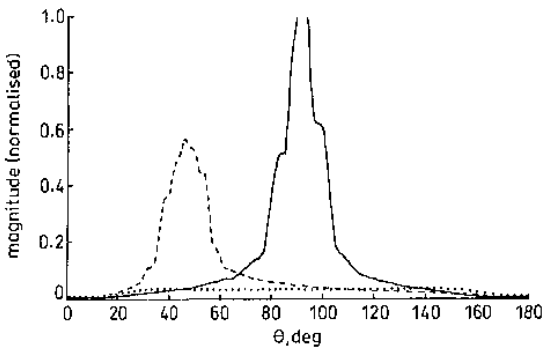

Fig. 10 *Polar plot of spectrum of directional surface*
—— rotated by $\phi = 0°$
····· rotated by $\phi = 90°$
--- rotated by $\phi = 45°$

assumption with the following argument: the surfaces are assumed to have random phase and be globally flat. The random phase spectrum condition allows us to invoke the central limit theorem, predicting that the surface height distribution will be Gaussian. Differentiation is a linear operation, therefore the distributions of the derivatives are also Gaussian. Since the surface is globally flat the derivative distributions will both have zero means. The rendering function is assumed to be linear, we would therefore expect the images to retain a Gaussian distribution that is symmetric about the intensity which corresponds to a horizontal facet. Assuming the reflectance function is identical for all the textures, the classes will share a common mean, which can be removed from the image giving a filter output with a mean value of zero.

The variance of the filter output is given by

$$\sigma^2 = \int_{\theta=0}^{2\pi} \int_{\omega=0}^{0.5 f_s} |G(\omega)|^2 \omega^2 |\sin\sigma|^2 |\cos(\theta - \tau)|^2 \times S(\omega, \theta - \phi)\, d\omega\, d\theta \tag{8}$$

For simplicity, assume that the surface spectrum is separable into radial and polar functions, $S_r(\omega)$ and $S_p(\theta)$

$$\sigma_g^2 = \int\int |\cos(\theta - \tau)|^2 S_p(\theta - \phi) \omega^2 |\sin\sigma|^2 \times S_r(\omega) |G(\omega)|^2\, d\omega\, d\theta \tag{9}$$

In the case of an isotropic surface, S_p is a constant and surface rotation will have no effect. If we approximate the polar plot of the directional surface oriented in direction γ, with an impulse function, $\delta(\theta - \gamma)$, i.e.

$$S_p(\theta) = \delta(\theta - \gamma - \phi) \tag{10}$$

Then eqn. 9 will be zero for all directions except $\theta = \gamma + \phi$, and

$$\sigma_g^2 = \int |\cos(\gamma + \phi - \tau)|^2 \omega^2 |\sin\sigma|^2 S_r(\omega) |G(\omega)|^2\, d\omega \tag{11}$$

We would therefore expect the standard deviation of the filtered image associated with the directional surface to vary with $|\cos(\gamma + \phi - \tau)|$ and that of the isotropic surface to be constant.

Before the filtered image can be passed to the classifier, it must be processed by a nonlinear function. In this paper we use a simple rectification operation. We assume the slope distribution to be Gaussian. Both Kube's model and the Gabor filters are linear operators, and we would therefore expect the signal prior to the nonlinearity to be Gaussian with standard deviation σ_g.

The rectification will cause the probability distribution to assume the form (eqn. 12)

$$p(f) = \frac{2}{\sqrt{2\pi}\sigma_g} \exp\left(\frac{-f^2}{2\sigma_g^2}\right) \tag{12}$$

where $f \geq 0$ is the value of the feature.

The mean of this distribution is approximately $0.6\sigma_g$. The mean of the feature derived from a unidirectional surface would therefore vary with the cosine of the angle between τ and θ.

$$\mu_F \propto |\cos(\gamma + \phi - \tau)| \tag{13}$$

Therefore, in the case of a directional surface, the first-order statistics of the IRIS feature are still functions of the angle of rotation, ϕ. This means that the classifier will be sensitive to surface rotation, and this may cause it to fail.

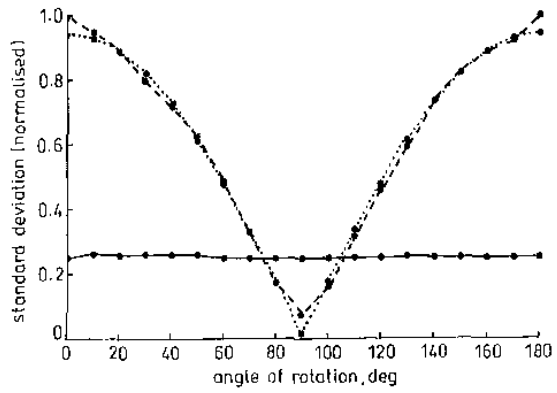

Fig. 11 *Standard deviations of Gabor filtered textures*
- •— rock
- –•·· sand
- ···•··· best fit

This effect is shown in Fig. 11 with the surfaces 'rock' and 'sand ripple', introduced in Figs. 2 and 3, respectively. The results for the sand ripple texture are well modelled with a best fit $|\cos(\gamma + \phi - \tau)|$ curve.

4 A rotation invariant scheme

The inspection task discussed earlier is designed to classify surfaces on the basis of their image texture. For a rotation invariant algorithm to be applied, the directional effect due to the illuminant must be removed. We now briefly consider several techniques.

4.1 Candidate techniques

4.1.1 Single image schemes: Chantler has used Kube's model as the basis for a scheme for the removal of directional effects associated with the illumination [9], while Pentland has used it as the basis for a surface shape estimator [15]. Both schemes estimate a quantity that is independent of the orientation of the surface relative to the illuminant, and which would therefore be a useful starting point for a rotation invariant classifier. However, both are derived from Kube's model, and share two significant weaknesses that stem from the linearisation inherent in Kube's model. First, the schemes are unable to estimate signal components perpendicular to the illuminant direction, and neighbouring estimates are noisy. Secondly, a linear model can be optimal for only a certain class of surfaces [16]. For a classification task, where several textures are present, the estimate will necessarily be suboptimal. In view of this we believe that neither of these schemes is sufficiently robust for use in a classifier.

4.1.2 Photometric schemes: In our hypothetical inspection task, we wish to classify surfaces. It would be desirable if we could do so on the properties of the surface, by eliminating the effects of illumination and reducing the problem to one of rotation invariance. Photometric techniques use several images of the same surface imaged under different illumination directions to form an estimate of surface shape. Unlike conventional shape from shading algorithms, the photometric problem is not underconstrained, and it is not necessary to use the smoothness constraints used in most single image shapes from shading algorithms. McGunnigle and Chantler [17] used photometric stereo to obtain surface derivatives of training surfaces which could then be rendered to produce training data appropriate to a surface illuminated from an arbitrary tilt angle. While this technique proved successful, it is not suitable for the problem of rotation invariance. Smith *et al.* [18] took the use of photometric stereo out of the context of training and also used it at the classification stage. Using the photometric estimate they were able to isolate the albedo component of texture from that due to the topography. In this paper we propose a scheme which combines photometric stereo with a modified rotation invariant algorithm to produce a classifier that is robust to surface rotation.

A photometric scheme will produce estimates of the surface derivatives which may be integrated to give a surface height map. A height map is initially attractive since it is a real scalar field that can be directly incorporated into an existing classifier. However, recovery of the height map from the surface derivatives is not trivial for rough surfaces; errors are cumulative and the derivative estimates may not be integrable. Furthermore, in spectral terms, integration represents the emphasis of low frequency components and the deemphasis of high frequency components. This is not desirable for segmentation.

Smith suggests the use of bump mapping, i.e. synthetically rendering the estimated derivatives from an angle that will emphasise discriminatory features. This will yield a scalar that can easily be accommodated into a texture analysis algorithm [18]. We note, however, that this still includes a directional filtering effect and the suppression of a significant amount of surface information. It would be preferable to classify surfaces on the surface derivatives themselves, rather than some scalar quantity derived from them.

4.2 An SRIS classifier

Classification on the basis of the surface derivatives presents two difficulties. First, the surface derivatives are a vector rather than a scalar quantity and the classifier must be applied to two random fields instead of one. Secondly, the components of this vector, i.e. the partial derivatives of the surface, contain a directionality which is an artefact of the process of differentiation. This directionality must be removed prior to classification.

In this paper frequency domain filtering (using the same set of filters used in the IRIS classifier) forms the first stage of feature extraction. The first problem is dealt with by treating the partial derivatives as the real and imaginary inputs to the Fourier transform. In this way the two fields can be filtered with relatively little computational overhead beyond that associated with the (scalar) IRIS filter [19, p.511]. (Note that the use of the complex FFT is purely for computational efficiency.)

The spectral representation of the derivative fields (eqns. 14 and 15) gives a clearer idea of the directionality associated with the derivative fields. The magnitude of these two expressions $M(\omega, \theta)$ (eqn. 16) is isotropic, i.e. the fields can be combined to form a quantity which does not contain any directionality apart from that inherent in the surface spectrum, $S(\omega, \theta)$.

$$P(\omega, \theta) = i\omega |\cos\theta| S(\omega, \theta) \tag{14}$$

$$Q(\omega, \theta) = i\omega |\sin\theta| S(\omega, \theta) \tag{15}$$

where $S(\omega, \theta)$ is the Fourier transform of $s(x, y)$, and $P(\omega, \theta)$ and $Q(\omega, \theta)$ are the Fourier transforms of $p(x, y)$ and $q(x, y)$, respectively.

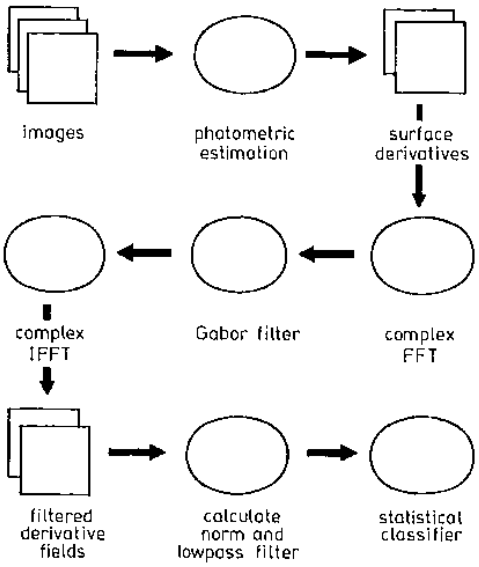

Fig. 12 *Outline of proposed SRIS scheme*

The magnitude of these two quantities is given by

$$M(\omega, \theta) = \sqrt{P(\omega, \theta)^2 + Q(\omega, \theta)^2} = i\omega S(\omega, \theta) \qquad (16)$$

The filtered derivative fields, $p_g(x, y)$ and $q_g(x, y)$, still contain a directionality due to differentiation which we wish to remove. Since we have already filtered the derivative fields we are only interested in the amount of energy contained in each frequency band. It is therefore permissible to use a nonlinear transformation at this stage without affecting the integrity of the spectral filtering. To remove the directional artefacts of the component derivative fields we take the norm of the vector field (eqn. 17). The resulting quantity is free of the directional filtering effect and is therefore a suitable basis for a rotation invariant algorithm. This operation is analogous to the rectification used in the IRIS classifier and, in combination with the same lowpass filtering used in the IRIS classifier, allows a statistical discriminant (of the same type and with the same number of features as that used in the IRIS classifier) to classify on the basis of the energy passed by each of the filters. The classifier is summarised in Fig. 12.

$$r_g(x, y) = \sqrt{p_g^2(x, y) + q_g^2(x, y)} \qquad (17)$$

Due to the use of the complex FFT there is little additional computational burden beyond that of the IRIS classifier. The additional operations are one image multiplication and one image addition for each feature image. Since these are both pixelwise operations they are relatively cheap compared to the forward and inverse Fourier transforms and lowpass filtering common to both the IRIS and SRIS classifiers

5 Results

5.1 Comparison of SRIS and IRIS classifiers

The new algorithm was applied to the original data set. At each rotation the surface montage was rendered using a Lambertian model and the surface derivatives estimated using a simple photometric system [16]. The estimates were passed to the modified classifier and the classification accuracy recorded. The boundaries of the classified surfaces under 0° and 45° of rotation are overlaid onto

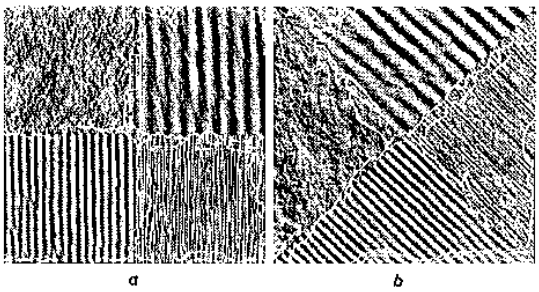

Fig. 13 *Classification of test montages by SRIS classifier*
a Under 0° rotation
b Under 45° rotation

the images in Fig. 13. The accuracy of the conventional image based (IRIS) classifier is shown for comparison in Fig. 14. The photometric classifier is clearly superior to the naïve IRIS classifier in terms of robustness to surface rotation. The photometric classifier is able to maintain a stable level of classification throughout the rotation range. This level is comparable with the best results obtained from the IRIS classifier at the original training angle. We believe this result justifies the extra expense of capturing the additional images.

5.2 Comparison with retrained IRIS classifier

As a second point of comparison we now compare the proposed scheme with an IRIS classifier that has been retrained at each rotation. Using a retrained classifier allows us to resolve the effect of changes in the characteristics of the training data from changes in the difficulty of the classification at a particular orientation.

The results show that the difficulty of classification based on the image textures varies significantly with orientation, (Fig. 15). In our montage, at 90° of rotation the prominent surface directionalities are perpendicular to the illuminant vector and are consequently heavily attenuated. The loss of this useful discriminatory information makes classification more difficult and is reflected in a significant increase in the level of misclassification.

Since the proposed scheme operates on the estimates of the surface derivatives, rather than image intensities, the information content of the data set is not only more consistent, but also richer, since directional components are not attenuated by the imaging process. In this way, the additional images required by the photometric system serve to make the classifier both more stable and more accurate.

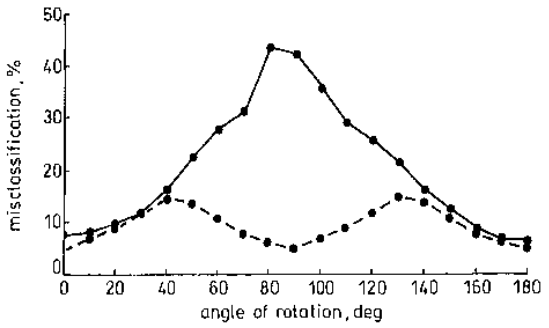

Fig. 14 *Variation of classification accuracy with angle of rotation*
—•— IRIS
– –•– – SRIS

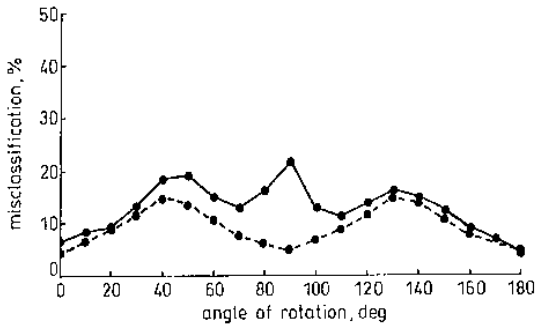

Fig. 15 *Comparison of SRIS scheme with retrained IRIS classifier*
– – • – – SRIS
• IRIS retrained

5.3 Robustness of the scheme

The proposed technique has, up to this point, been applied to test data simulated under idealised conditions. We now test the robustness of the algorithm to additive noise. Comparison is made with the IRIS classifier which has been trained at $\phi = 0°$ and both algorithms are applied to the data set that has been rotated 45°. We assume the noise is white (and consequently isotropic) and Gaussian and that it is present in both training and test images. Various realisations of the process are added to the photometric recovery images. Since the variance of the images varies with the surface orientation, the signal to residue ratio associated with noise processes of a given variance will vary. The ratio quoted is the average of the values associated with each three photometric recovery images.

In fact both the IRIS and SRIS classifiers are surprisingly tolerant to noise (Fig. 16). This is largely due to the nature of the textures: even with at signal to noise ratio of -5.85 dB it is possible to visually discriminate between the textures (Fig. 17). Comparison with the IRIS classifier shows the photometric scheme to be less robust to noise. Below a signal to residue ratio of 2 dB the performance of the proposed scheme is inferior to the IRIS classifier at 45°. Furthermore the accuracy of the SRIS classifier collapses catastrophically, compared with the much more graceful degradation of the IRIS classifier. Although these results are highly dependent on the data set it does appear that the SRIS classifier is less robust to noise and is not an effective approach in noisy applications.

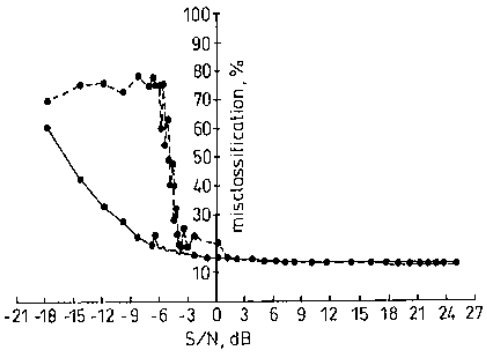

Fig. 16 *Effect of additive image noise*
— • — IRIS
• SRIS

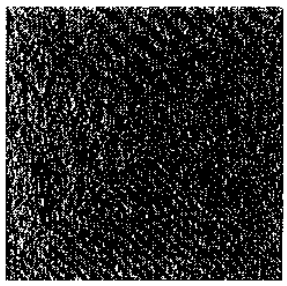

Fig. 17 *Test montage corrupted by white noise*
SNR = -5.85 dB

6 Conclusions

In this paper we have highlighted the difference between rotation of a textured image and rotation of the textured surface from which it is formed. Many visual textures are formed by the interaction of light with a rough surface, and in many applications the classifier must cope with rotated surfaces. We have shown that, when directional surfaces are present an image rotation insensitive approach can fail badly under surface rotation. This problem is most significant when the illuminant vector is nearly horizontal, unfortunately this is also the condition under which the surface texture becomes most apparent. We therefore argue that many of the rotation invariant algorithms presented in the literature are inappropriate for classification tasks involving surface rotation.

We have proposed a scheme to overcome this problem in the context of an inspection task where illumination conditions can be controlled. The scheme uses photometric stereo to estimate the surface derivative fields. By using characteristics of the surface, rather than of the image, it is possible to apply a variant of a traditional rotation invariant scheme. However, the estimated partial derivative fields represent a two-dimensional vector quantity, rather than the scalar field associated with most texture analysis. By treating the two scalar, partial derivative fields as one complex field it is possible to efficiently implement a frequency domain filter without additional computation during the forward and backward transforms. The norm of the filtered derivative fields gives a feature that is free from the directional filtering effect inherent to differentiation.

The proposed scheme's performance was tested using simulation. Where the noise level is moderate the SRIS scheme was found to be more robust to surface rotation than the IRIS technique. Furthermore, it was shown that the scheme is able to utilise information from the photometric recovery images which can offer improvements in classification performance over that of the IRIS classifier.

7 References

1 HARALICK, R.M., SHANMUGAM, K., and DINSTEIN, I.: 'Textural features for image classification', *IEEE Trans. Syst. Man Cybern.*, 1973, **SMC-3**, pp. 611–621
2 SUN, C., and WEE, W.G.: 'Neighboring gray level dependence matrix for texture classification', *CVGIP*, 1983, **23**, pp. 341–352
3 KASHYAP, R.L., and KHOTANZAD, A.: 'A model-based method for rotation invariant texture classification', *IEEE Trans. Pattern Anal. Mach. Intell.*, 1986, **PAMI-8**, pp. 472–481
4 MAO, J., and JAIN, A.K.: 'Texture classification and segmentation using multiresolution simultaneous autoregressive models', *Pattern Recognit.*, 1992, **25**, (2), pp. 173–188
5 COHEN, F.S., FAN, Z., and PATEL, M.A.S.: 'Classification of rotated and scaled textured images using Gaussian Markov field models', *IEEE Trans. Pattern Anal. Mach. Intell.*, 1991, **13**, pp. 192–202

6 PORTER, R., and CANAGARAJAH, N.: 'Robust rotation invariant texture classification: wavelet, Gabor and GMRF based schemes', *IEE Proc. Vis. Image Signal Process.*, 1997, **144**, (3), pp. 180–188
7 CHEN, J.L., and KUNDU, A.: 'Rotation and gray scale transform invariant texture identification using wavelet decomposition and hidden Markov model', *IEEE Trans. Pattern Anal. Mach. Intell.*, 1994, **16**, (2), pp. 208–214
8 HALEY, G.M., and MANJUNATH, B.S.: 'Rotation-invariant texture classification using modified Gabor filters'. Proceedings of IEEE conference on *Image Processing*, 1996, **1**, pp. 262–265
9 CHANTLER, M.J.: 'The effect of illuminant direction on texture classification'. PhD Thesis, Heriot Watt University, Edinburgh, 1994
10 CHOE, Y., and KASHYAP, R.L.: '3-D shape from a shaded and textural surface image', *IEEE Trans. Pattern Anal. Mach. Intell.*, 1991, **13**, pp. 907–919
11 FUKUNGA, K.: 'Introduction to statistical pattern recognition' (Academic Press, London, 2nd edn.)
12 OGILVY, J.A.: 'Theory of wave scattering from random rough surfaces', (Adam Hilger, Bristol)
13 KUBE, P., and PENTLAND, A.: 'On the imaging of fractal surfaces', *IEEE Trans. Pattern Anal. Mach. Intell.*, 1988, **PAMI 10**, (5), pp. 704–707
14 CHANTLER, M.J., RUSSELL, G.T., and LINNETT, L.M.: 'Illumination: a directional filter of texture?'. Proceedings of BMVC94, 1994, **2**, pp. 449–459
15 PENTLAND, A.: 'Linear shape from shading', *Int. J. Comput. Vis.*, 1990, **4**, pp. 153–162
16 McGUNNIGLE, G.: 'The classification of textured surfaces under varying illuminant direction'. PhD thesis, Heriot Watt University, Edinburgh, 1998
17 McGUNNIGLE, G., and CHANTLER, M.J.: 'A model-based technique for the classification of textured surfaces with illuminant direction invariance'. BMVC97, Proceedings of the British *Machine Vision* conference, 1997, **2**, pp. 470–480
18 SMITH, M.L., HILL, T., and SMITH, G.: 'Surface texture analysis based upon the visually acquired perturbation of surface normals', *Image Vis. Comput.*, 1997, **15**, pp. 949–955
19 PRESS, W.H., TEUKOLSKY, S.A., VETTERLING, W.T., and FLANNERY, B.P.: 'Numerical recipes in C, the art of scientific computing' (Cambridge University Press, 1995, 2nd edn.)