A Comparison of Three Rough Surface Classifiers

G. McGunnigle M.J. Chantler

G. McGunnigle and M.J. Chantler are with the Department of Computing and Electrical Engineering, Heriot Watt University, Riccarton, Edinburgh EH14 4AS, United Kingdom. E-mail: gmg@cee.hw.ac.uk

List of Figures

Abstract

In this paper texture analysis techniques are used to segment rough surfaces into regions of homogeneous texture. The performance of three rough surface classifiers was assessed and compared. The classifiers differ in their discrimination as well as their input and computational requirements. Simulation and experiment were used to identify the limitations of the classifiers and to identify which classifier is best suited to a particular task. A series of guidelines for the choice of classifier are presented and justified.

Keywords

Surface classification, texture analysis, rotation invariance, photometric stereo, Gabor filters

I. Introduction

The visual texture of a surface is caused by light interacting with variations in the surface's reflection characteristics or with the surface topography. This paper is limited to the second source of texture and deals with the classification of rough surfaces on the basis of their image texture. Although many investigations have used this form of texture *implicitly*—the majority of Brodatz textures [1] contain at least a component due to surface topography—little work has been carried out on the phenomena associated with this group. One characteristic of rough surface textures is that the appearance of the surface is a function of the illuminant direction as well as of the surface topography [2] [3] [4] [5]. For instance, rotation of a directional surface is not equivalent to rotation of its image, [6], Figure 1. A surface rotation invariant classifier must take this effect into account.

Texture analysis is an established field within computer vision. Because of the abundance of textures in unconstrained environments and because of the importance of texture as a stimulus to the human visual system, most research takes the image as the starting point. The development of rotation invariant techniques for texture analysis has been extensively researched [7] [8] [9] [10] [11] [12] [13]. However, these algorithms are invariant to rotation of the *image*. This is not appropriate to rough surfaces—the visual texture on the right of Figure 1 is completely different from the texture on the left of the figure. This paper will assess when it is possible to classify using the image and when it is necessary to infer surface properties.

The effect of directional illumination on classification can be reduced by decreasing the (slant) angle between the light source and the vertical—though of course this suppresses the visible texture. An alternative approach is to classify using the properties of the surface rather than those of the image. Photometric stereo allows the estimation of surface derivatives using several images of the same scene captured under different illumination conditions. Critically, it does not need the smoothness constraint required by most single-image shape from shading algorithms. Figure 2 shows a series of scatter plots of the estimated surface derivatives, p and q, corresponding to the image series (Figure 1). The joint distribution of p and q is invariant, apart from a rotation. This makes the surface derivatives an attractive feature for classification.

This paper compares the ability of three test classifiers—IRIS [11], Point [2] and SiRIS [6]—to distinguish and segment Gaussian textured surfaces that obey Lambert's law. The IRIS and SiRIS classifiers have never before been tested on real, surface rotation data and the Point classifier has only been previously tested on a limited data set. This paper is also novel in its approach to assessment: we constrain the classification tasks so that the classifiers must discriminate between similar textures. In this way, we can identify the types of surface to which each classifier is suited, though this also results in relatively high error rates.

The techniques are tested in an inspection task where imaging conditions are controlled: we do not try to mimic the versatility of the human visual system. In many inspection tasks the classifier deals with surfaces that have been formed by similar physical processes—that is they differ in the degree, rather than the type, of processing. The experimental approach of this paper is to test the ability of the classifiers to distinguish similar surfaces. Initially the investigation uses simulation with physically based surface models to assess the technique. By varying the parameters of the model it is possible to control the similarity of the surfaces. This allows the relative strengths of the classifiers to be assessed methodically. The natural textures used in this paper are grouped in montages according to their underlying physical cause. We have tried to use sets of surfaces that are similar, but which show a gradation in appearance. The natural surfaces are used to assess the relative performance of the classifiers as well as their robustness to surface rotation.

We have found that the Point classifier is the least able to discriminate between surfaces and is the most susceptible to low frequency trends. The SiRIS classifier performs better than IRIS—except when the illuminant can be aligned to accentuate differences in surface directionality. Most significantly we have confirmed that SiRIS is robust to surface rotation (in the plane); and that IRIS is not.

Although the Point classifier is the worst classifier tested we do not rule it out for applications: its simplicity makes it attractive where fast detection, or a fast, coarse classification, is required. For isotropic surfaces SIRIS improves on IRIS with little increase in computation—though whether the improvement justifies the need for more images depends on the application. For directional surfaces that are always presented at the same orientation and when the lighting can be optimised, we recommend the IRIS classifier. If the orientation of a directional surface cannot be predicted—we conclude that IRIS is insufficient and a photometric approach such as SiRIS is necessary.

II. Test Classifiers

Tuceryan and Jain identify four general approaches to texture classification: statistical approaches, geometric methods (including structural techniques), model-based techniques and signal processing methods [14]. The classifiers we consider fall into the signal processing category—filtering followed by an energy measure. Recently several authors have used alternatives to energy measures. Pietikäinen et al. first of all compare feature distributions to select the most likely class, then apply the k nearest neighbour algorithm to assign the final classification [15]. The Point classifier omits the filtering operation—however in all other respects it is equivalent to the other classifiers. This class of technique is derived from classical detectors and a discriminant is used to set a series of thresholds which determine the mapping from feature value to estimated class. This may be implemented in several ways, e.g. neural nets [16], nearest neighbour classifiers [18], or statistical classifiers [17]. We used a simple Bayes Linear classifier, see [28], with all the classifiers tested. We chose a linear statistical classifier because of its popularity in the literature. This paper deals with the robustness of features to rotation: the performance of different discrimants will vary, however they are all beset by the same problem of feature variation, and the differences in their performance will, in general, be small relative to the variations caused by surface rotation.

In [2] we proposed a classifier (Point) that used photometric stereo to estimate the surface derivatives then extracted features from the point statistics of the derivative fields. In this classifier feature extraction corresponds to measuring the local distribution of the surface derivatives, p and q, and estimating its major and minor axes. The features encapsulate surface roughness and directionality and are rotation invariant. Subsequently, we have concluded that the sum of the squares of the surface derivatives, although a less powerful feature, is more robust and better suited for classification: this is the form used in this paper. This is similar to the approach pursued by Smith [25], where he extracted a series of descriptive features from the point statistics of derivative fields. In [6] a Surface Rotation InSensitive (SiRIS) classifier was proposed that used isotropic Gabor filters and was *insensitive* to surface directionality and therefore robust to surface rotation. The algorithm first decomposes the surface derivative fields into a series of bandlimited fields, then calculates the sum of the squares of the filtered surface derivatives for each band. The third technique investigated is the Image Rotation InSensitive (IRIS) classifier. This is based on the scheme proposed in [11], though in this paper it is used to classify on a pixel by pixel basis rather than to classify image blocks. SiRIS is almost identical to IRIS except that IRIS operates on image data rather than inferred surface characteristics. The three classifiers are summarised in Figure 3.

All of the classifiers considered ignore phase and discriminate on a sampling of the power spectrum of either the image or the surface derivatives. A stationary random process may be completely characterised by its power spectrum if the process is Gaussian. The assumption that the surface height distributions are Gaussian is common in the rough surface literature. Experimental results suggest that, although not universal, the Gaussian assumption is valid for many natural surfaces [22]. Since differentiation is a linear operation, the assumption also holds for surface derivatives. Furthermore, if the surface facets are not too steep, Lambertian reflection may be modelled as a linear operator [23], and the Gaussian assumption can be extended to the distribution of image intensities. In this paper, surfaces are assumed to be characterised solely by their power spectrum.

The classifiers differ in their requirements for data and computation. The IRIS scheme, which is a conventional image-based classifier, needs only one image. The Point and SiRIS classifiers rely on photometric stereo and require a minimum of three images for a Lambertian surface. However in this paper we use four images because we have found that a fourth image makes the estimator more robust to shadows. The Point classifier requires least computation. SiRIS and IRIS require similar amounts of calculation. Although SiRIS applies each filter to two images (or derivative fields), whereas IRIS applies each filter to only one, the image pair can be treated as being a single, complex image and filtered simultaneously, see [19] P.511 for explanation. In this way there is no substantial overhead for filtering pairs of surface derivatives rather than single images. This point is

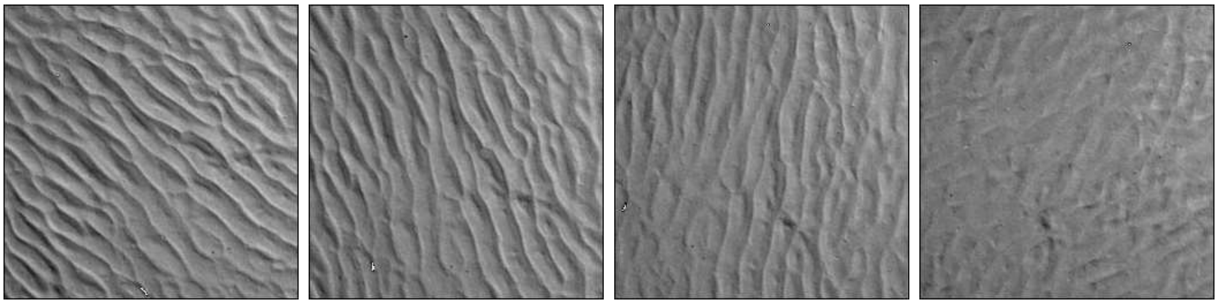

Fig. 1. Directional surface rotated clockwise in 30° increments.

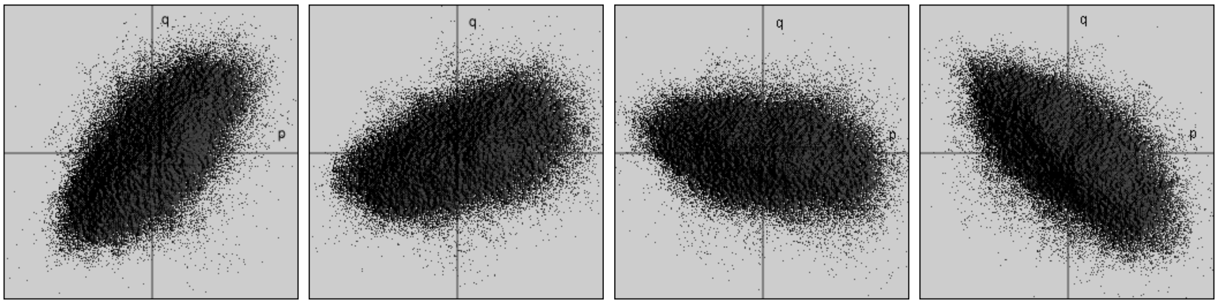

Fig. 2. Distribution of facet slope for the rotated directional surface. (Both axes range from -0.25 to 0.25).

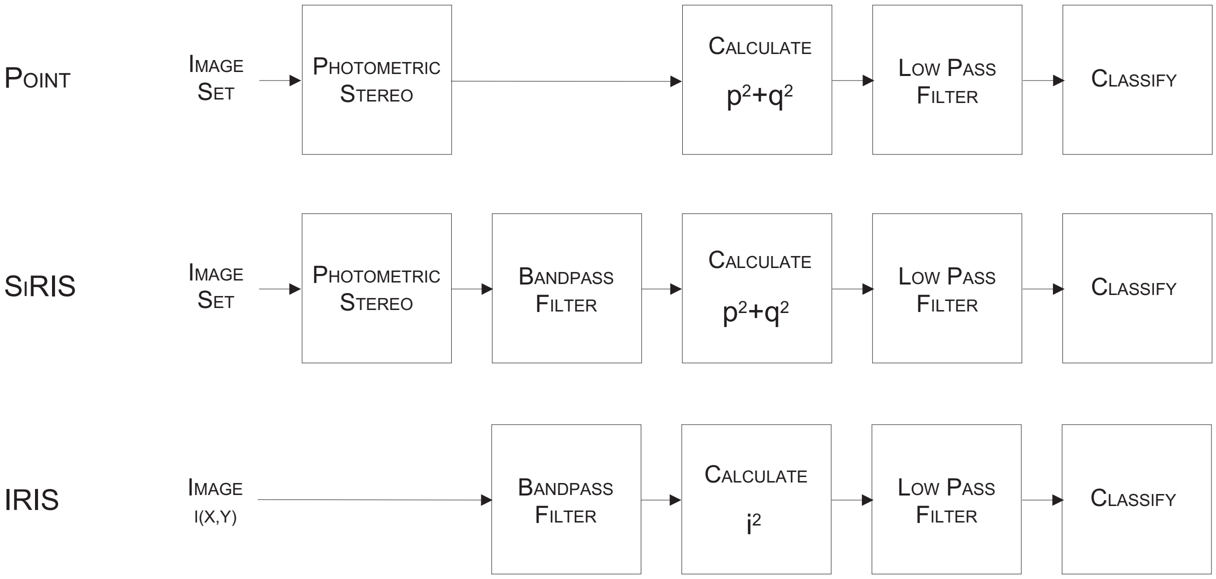

Fig. 3. The three test classifiers.

confirmed by comparing the CPU times of the different algorithms, Table I. Point is the fastest by a large margin and IRIS is only slightly faster than SiRIS. Given the relative costs of the classifiers it is desirable to identify when the extra expenditure is necessary or useful.

III. Simulations with synthetic surfaces

The classifiers' ability to distinguish two classes of surface was assessed. Simulation using parametric surface models was used to vary the similarity of the test surfaces—and the difficulty of discrimination—in a controlled manner.

A. Method

To allow comparison, all the classifiers used the same training data and the same Gaussian low pass filter (with a standard deviation of 4 cycles per image). The IRIS and SiRIS classifiers also shared the same sampling of the frequency spectrum: Gabor filters with centre frequencies of 32, 64, 96 and 128 cycles per image (c/i) and a Gaussian low pass filter, all with a standard deviation of 16 c/i. A Bayes linear discriminant was used to make the final classification for each algorithm. The classifiers were trained on one sixteenth of the test image. Each montage is 256×256 pixels and consists of two textures. The image textures were formed by rendering realisations of stochastic surface models using Lambert's law. The synthetic surfaces were illuminated from a slant of $50°$, and tilt angles of $0°$, $90°$, $180°$ and $270°$. As the surfaces become more similar, the error rate approaches 50%. In fact it rarely reaches this value because information specific to each realisation is implicitly incorporated into the discrimination rule.

The surface models are parametric expressions of the power spectrum of the surface height field. Phase spectra are random—different realisations of a model are obtained using different phase spectra. Two isotropic surface models were used: *Sayles* [20] and *Mulvaney* [21]. The two dimensional expressions are shown in Equations 1 and 2, respectively. A directional model, *Ogilvy* [22], is shown in Equation 3.

$$S_{say}(\omega) = \frac{k_{say}}{\omega^3} \tag{1}$$

$$S_{mal}(\omega) = k_{mal}\left(\frac{\omega^2}{\omega_c^2} + 1\right)^{-\frac{3}{2}} \tag{2}$$

$$S_{ogl}(u, v) = \frac{k_{ogl}}{(u^2 + {\omega_a}^2)(v^2 + {\omega_b}^2)} \tag{3}$$

where

$S(\omega)$	is the two dimensional power spectrum
ω	is the radial frequency
ω_c	is the cut-off frequency
k_{say}, k_{mal} and k_{ogl}	are constants
u and v	are the cartesian frequency coordinates
ω_a and ω_b	are the cut-off frequencies in the x and y directions.

The *Sayles* model is fractal and is parameterised by the factor k_{say}, or equivalently by the root mean square (rms) slope. This parameter is varied in Simulation 1. The classifiers must make a series of discriminations between two fractal surfaces: the first has an rms slope of 0.25, the rms slope of the other is varied between 0.21 and 0.29. Figure 4 shows the effect of varying rms slope from 0 to 0.5. The second simulation repeats the first, except that a *Mulvaney* surface, with a cut-off frequency of 32 c/i, is used. The accuracy of classification for Simulations 1 and 2 is shown in Figures 7 and 8 respectively.

The *Mulvaney* model has an additional parameter that controls the frequency at which the transition from white noise to fractal behaviour occurs: when the cut-off frequency is low the model approximates the *Sayles* model. In Simulation 3 the first surface has a cut-off frequency of 16 c/i and the cut-off frequency of the second is varied from 1 to 32 c/i. Unlike the previous experiment, the test surfaces now differ in the shape as well as the power of the spectrum. As we change the cut-off frequency we rescale the power spectrum so that the amplitude of the high frequencies remains unchanged. This is designed to simulate a wear process which affects the high amplitude, low frequency, components first. The visual effect is shown in Figure 5. The accuracy of classification is shown in Figure 9, and an example segmentation for each classifier is shown in Figure 10.

The *Ogilvy* model has three parameters: the scaling parameter, k_{ogl}, and the vertical cut-off frequency, ω_b, are held constant for both surfaces (ω_b=4); the horizontal cut-off frequency, ω_a is set to 16 c/i for one surface and varied from 1 to 32 c/i for the second. The visual effect is shown in Figure 6. Since the appearance of a directional surface varies with the direction of illumination, IRIS is tested for two cases: when the illuminant is perpendicular to the surface grain (IRIS0) and when it is parallel (IRIS90). The photometric classifiers use illumination from both of these directions, so only one set of results is needed for each of these classifiers. The effect on the classifiers is shown in Figure 11

B. Simulation Results and Discussion

The Point classifier is significantly less accurate than the SiRIS classifier in Simulation 1, Figure 7. This is surprising since it directly and parsimoniously measures the experimental variable. However, when the experiment was repeated with the *Mulvaney* surface (Simulation 2) it performs well relative to the others. This suggests that the Point classifier is more sensitive to low frequency trends than the spectral classifiers which isolate—and can therefore ignore—low frequency signal components.

The Point classifier also performs more poorly than the other classifiers in Simulations 3 and 4 where roughness is varied indirectly; though it does outperform the IRIS classifier in Simulation 4 when the surface grain is parallel with the illuminant, Figure 11. Despite this poor performance, measured in absolute terms it is able to discriminate between similar surfaces and may be sufficient in many applications.

SiRIS classifies using the filtered surface derivatives: IRIS uses the filtered image intensity. Nonetheless, they can be usefully compared. The reflectance function may be approximated as being linear: that is the intensity of a facet can be treated as the weighted sum of its partial derivatives. This simplifies comparison of the two techniques.

In the case of isotropic textures (Simulations 1-3) SiRIS is, in most cases, superior (Figures 7-9). For an isotropic surface the linear approximation allows the image to be treated as a single scaled and rotated derivative field. In this way the IRIS classifier discriminates on the basis of the square of one (scaled) derivative field. The SiRIS classifier discriminates on the sum of the squares of two derivative fields. Since the surface is isotropic, the derivative fields are uncorrelated and have identical distributions. The sum

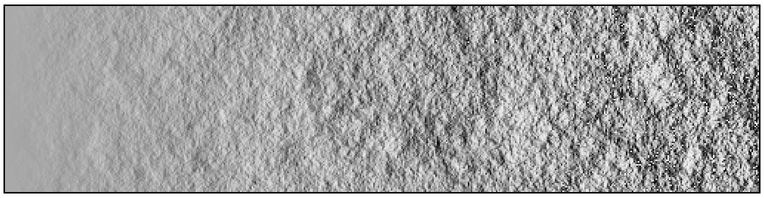

Fig. 4. Sayles surface: the effect of varying surface roughness.

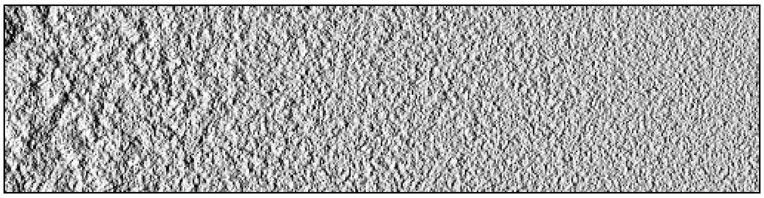

Fig. 5. *Mulvaney* surface: the effect of varying cut-off frequency.

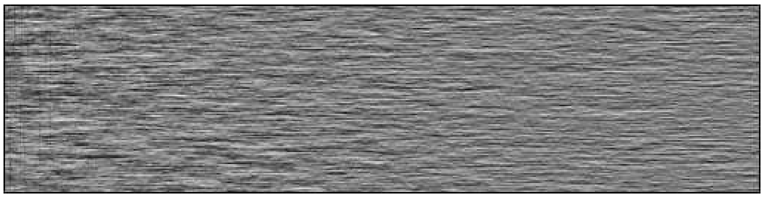

Fig. 6. *Ogilvy* surface: the effect of varying horizontal cut-off frequency.

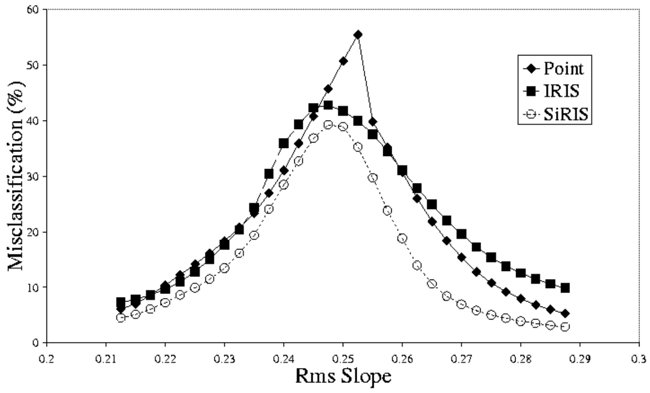

Fig. 7. Simulation 1: Accuracy of classifiers with scaled *Sayles* surfaces

of the squared fields will have twice the mean but only $\sqrt{2}$ times the standard deviation of a single squared field. A feature based on two fields will have less overlap between distributions of different classes, and therefore less misclassification.

In Simulation 3 the superiority of SiRIS is most pronounced where the surfaces differ at low frequencies. A given low pass filter will be less effective at reducing the variance of a low frequency band than of a high frequency band. It follows that any scheme that increases the mean relative to the standard deviation will make its most significant improvements at low frequencies. This explanation is consistent with the simulations.

For the directional surface (Simulation 4) the performance of the IRIS classifier depends on the direction of the illuminant relative to the orientation of the surface, Figure 11. This is due to the directional filtering effect of the illuminant demonstrated in Figure 1. If the surface grain is parallel to the light source (IRIS90) IRIS performs significantly worse than both the Point and SiRIS classifiers. If the surface grain is perpendicular to the illuminant (IRIS0) it is more effective than either. Parallel to the grain the illuminant suppresses the differences between the surfaces; perpendicular it accentuates them and makes discrimination easier. In the case of *Ogilvy* simulations the differences are predominantly in the x-direction and are concentrated in the p-derivative field—the q-field effectively adds noise to the classification. When the surface is oriented in the optimum direction the directional effect of the illuminant attenuates this noise—making IRIS more accurate.

The simulations lead us to the following conclusions: first the Point classifier has least discrimination and is the most vulnerable to low frequency trends. Secondly, SiRIS is consistently more accurate than IRIS, except where the test surfaces differ in one direction and the illuminant can be constrained to accentuate these differences. Finally, the IRIS classifier is sensitive to the relative orientation of the illuminant and the surface directionality.

IV. Experiments with natural surfaces

A. *The Effect of non-Lambertian Reflection*

Both the Point and SiRIS classifiers use photometric stereo to recover the surface derivatives. For classification the estimates need not be accurate, but they must be consistent.

In this section we test the consistency of the estimate for groups of surfaces that violate, to varying degrees, the assumptions of this paper.

The sample is held stationary and two sets of photometric images are captured; one with $\tau = 0°,90°,180°$ and $270°$, and one with $\tau = 60°,150°,240°$ and $330°$. The two estimates of the p-derivative were compared and the standard deviation of the difference measured. We used three groups of surfaces: Fractures—these fit our assumptions and are our control group, though the roughest exhibits obvious shadowing. The repeating primitives group have a specular component in addition to the diffuse component of reflection, and exhibit significant shadowing. The vertically milled surfaces are metal and highly specular, unlike the other samples these are lit from a slant angle of 75° to reduce the amount of specular reflection reaching the camera.

The consistency of the estimate, Table II, is clearly related to the group: the control group—which fit our assumptions—give the most consistent results; the repeating primitive group show some degradation; the vertical milled surfaces cannot be consistently estimated with this photometric algorithm and we will not pursue any further experiments with these surfaces.

B. Classification Experiments

The classification experiments with natural surfaces had two aims: to assess the classifiers' discrimination and to test the robustness of the classifiers to surface rotation.

The classifiers used for the experiments on natural surfaces were identical to those used in the simulations except that the IRIS and SiRIS classifiers sample the spectrum at 0, 64, 128, 192 and 256 (c/i) with a standard deviation of 32 c/i. An extra post-processing stage was also added, the classified image was filtered with a 10×10 pixel mode filter. The classifiers were trained on one sixteenth of the image, i.e. a 64×64 pixel block from the centres of each test sample. In the case of rotation experiments the classifiers were trained prior to rotation.

Photometric image sets were captured by illuminating the surfaces from a slant angle of 50°, and tilt angles of 45°, 135°, 225° and 315°. The first four groups of natural surfaces are made of plaster and are near-Lambertian. This simplifies photometric estimation, however, photometric stereo is not limited to this case and has been applied to a much wider range of surface reflectance functions both with, and without, prior calibration, [26] [27].

The test surfaces (Figure 13) are grouped into one of five (512×512) montages, depending on the physical process that formed them. *Fracture* surfaces were formed by impact fractures of solid blocks of plaster. They vary from the relatively smooth, ceramic-like fracture of *Fracture 1* to the much rougher, more fractal, *Fracture 4* surface. *Deposit* surfaces were formed by the deposition of plaster powder on a plaster flat. These surfaces differ in the amount of powder deposited on the surface. *Ground* surfaces were formed by grinding plaster flats. *Ground 1* was prepared by grinding in two orthogonal directions; the other surfaces in this group are unidirectional and vary because of the different length of grinding stroke used in their preparation. *Ripple* surfaces were formed by wave action and differ in directionality and frequency.

In addition to these groups we have also included a group of surfaces that are outwith the assumptions of the paper. The *Repeating Primitive* surfaces are composed of many, randomly placed and oriented, primitives. This group violate three assumptions: they have specular reflection; they are phase rich and exhibit significant shadowing.

The natural surfaces are used in five experiments. The first experiment assesses the ability of the classifiers to discriminate between similar surfaces, i.e. we apply the classifiers to montages of the samples from each group, Table III. The remaining experiments test the robustness of the classifiers to surface rotation. We could not include the Fracture 1 sample in this experiment: the physical sample is too small to completely fill the image when the sample is rotated. Rather than use different imaging conditions from the rest of the samples we decided not to use this surface for rotation experiments.

We combine *Fractures 2* and *4* and *Deposits 1* and *4* to form the *Iso* montage. The classifiers are trained at $0°$ of rotation and tested on surfaces rotated in $30°$ increments. At each rotation a photometric image set is captured.

C. Experimental Results and Discussion

In Experiment 1 the ability of the classifiers to discriminate between closely related surfaces was tested (Table III). The segmentations of the *Deposit* montage are shown in Figure 14. For all montages the SiRIS classifier was the most effective and the Point classifier was the worst. However, except for the *Ripple* montage, the Point classifier is performing a degree of discrimination; this may be adequate for some applications. Both the simulations and the experimental results indicate that SiRIS is a more robust classifier than IRIS, however, the experiments show a greater advantage than we would have predicted from the simulations.

Experiment 2 tests the classifiers' robustness to the rotation of isotropic surfaces. Unsurprisingly, all the classifiers are unaffected (Figure 15). The relative accuracies are similar to those found for the isotropic montages in Experiment 1.

Experiments 3 and 4 (Figures 16 and 17) test the classifiers' robustness to rotation of directional surfaces. In both experiments the IRIS classifier fails catastrophically as the surface is rotated. This is due to the directional filtering effect of illumination shown in Figure 1 and Simulation 4. This is a limitation of single image classifiers. As in Experiment 1 the Point classifier is completely unable to discriminate any of the *Ripple* textures. It *is* able to discriminate between the Ground surfaces—though it does show a degree of sensitivity to rotation: classification is most accurate when the surface directionality is aligned with a pair of illuminants. The SiRIS classifier gives the lowest misclassification and maintains a low level in both experiments.

In our final experiment we tested the classifiers on a series of phase rich surfaces that showed specular reflection and shadowing. Our aim was to assess how breaking our assumptions affected the photometric classifiers. In the event both the SiRIS and Point classifiers performed consistently, with SiRIS giving the better classification, Figure 18. IRIS, as in Experiment 2, is more accurate than Point but less accurate than SiRIS. However the main point is that the photometric classifiers do not seem to be adversely affected by either the shadowing or specularity present in this data set.

It is useful to reconsider the simulations in the context of these results. The simulations led us to three conclusions. First, that *Point* had least ability to discriminate surfaces than the other classifiers. Secondly, SiRIS is more accurate than IRIS except when the surface directionality is aligned with the lighting vector. Thirdly, IRIS is sensitive to the realtive orientation of the lighting and surface directionality. The first and third conclusions are fully vindicated by the experiments with real surfaces. The second conclusion is partially vindicated—we found that even when the surface was aligned with the lighting SiRIS was still more accurate. One explanation is that the real data is more noisy than the synthetic data and that SiRIS, using more images, is more robust to that noise.

V. Conclusions

In this paper three techniques that classify rough surfaces by their appearance were compared. Simulations and experiments on natural surfaces indicate that the SiRIS classifier is the most effective. In most cases the IRIS classifier is slightly poorer; however, if the test surfaces are directional, and their orientation is unpredictable, the IRIS classifier can fail catastrophically. The Point technique has the least ability to distinguish surfaces. It is able, however, to discriminate to some degree all the surfaces except those in the *Ripple* montage.

The choice of classifier is specific to the application. Where a rapid, rough classification is required, the computational simplicity of the Point classifier makes it an attractive option. If finer discrimination is required then one of the spectral classifiers should be used. We recommend IRIS if surfaces are presented at a known orientation and the illuminant direction can be optimised. If the classifier must deal with directional surfaces that are presented at arbitrary orientations then the IRIS classifier is inadequate and SiRIS should be used. If the surfaces are isotropic the choice is less clear: SiRIS does give more accurate classification than IRIS with no significant increase in computation, however, it does require a photometric set. Whether the improvement in accuracy outweighs the extra imaging requirements will depend on the application.

References

[1] P. Brodatz *Textures: a phtographic album for artists and designers*, Dover, New York, 1966.

[2] G. McGunnigle, M.J. Chantler, *Rough surface classification using point statistics from photometric stereo*, Pattern Recognition Letters 21 (2000) 593-604.

[3] K.J. Dana, B. van Ginneken, S.K. Nayar, J.J. Koenderink *Reflectance and texture of real world surfaces.* ACM Transactions on Graphics, Vol. 18, No.1, pp.1-34, January 1999

[4] P.-h. Suen, G. Healey *The Analysis and Recognition of Real-World Textures in Three Dimensions.* IEEE Transactions on Pattern Analysis and Machine Intelligence, May 2000, Vol.22 No.5 pp.491-503

[5] K.J. Dana and S.K. Nayar *Correlation Model for 3D Textures* , ICCV '99, pp. 1061-1067, September 1999.

[6] G. McGunnigle, M.J. Chantler, *Rotation invariant classification of rough surfaces.* IEE Proceedings Vision, Image and Signal Processing, Vol. 146, Number 6, December 1999.

[7] C. Sun, W.G. Wee, *Neighbouring gray level dependence matrix for texture classification*, CVGIP Vol.23, 1983, pp.341-352

[8] R.L. Kashyap, A. Khotanzad *A model-base method for rotation invariant texture classification*, PAMI Vol. 8, July 1986, pp.472-481.

[9] J. Mao, A.K. Jain *Texture classification and segmentation using multiresolution simulataneous autoregressive models*, Pattern Recognition Vol.25, No.2, 1992 pp.173-188.

[10] F.S. Cohen, Z. Fan, M.A. Patel *Classification of rotated and scaled textured images using Gaussian Markov field models.* PAMI Vol.13, Feb. 1991, pp.192-202.

[11] R. Porter and N. Canagarajah *Robust rotation invariant texture classification: wavelet, Gaborand GMRF based schemes*, IEE Proc. Vis. Image Signal Processing. Vol. 144, No.3, June 1997.

[12] J.L. Chen, A. Kundu *Rotation and Gray Scale Transform Invariant Texture Identification Using Wavelet Decomposition and Hidden Markov Model*, IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol.16, No.2, February 1994 .

[13] G.M. Haley, B.S. Manjunath *Rotation-invariant texture classification using modeified Gabor filters*, IEEE Conference on Image Processing, 1996, Vol.1, pp.262-265.

[14] M. Tuceryan, A.K. Jain *Texture Analysis*, Handbook of Pattern Recognition and Computer Vision, pp.235-276. Eds. C.H. Chen, L.F. Pau and P.S.P. Wang World Scientific Publishing Company

[15] M. Pietikäinen, T. Ojala and Z. Xu *Rotation-invariant texture classification using feature distributions*, Pattern Recognition, Vol.33, No.1, January 2000, pp.33-52.

[16] H. Greenspan, R. Goodman, R. Chellappa, and C.H. Anderson Learning Texture-Discrimination Rules in a Multiresolution System PAMI Vol. 16, 1994, No.9, September, pp.894-901

[17] Unser, M. Texture Classification and Segmentation Using Wavelet Frames IP Vol.4, 1995 No.11 November pp.1549-1560

[18] P.P. Ohanian and R.C. Dubes Performance Evaluation for Four Classes of Textural Features", PR Vol.25 1992 pp.819-833

[19] W.H. Press, S.A. Teukolsky, W.T. Vetterling and B.P. Flannery *Numerical Recipes in C The Art of Scientific Computing* Second Edition, Cambridge University Press.

[20] R.S. Sayles, T.R. Thomas, *Surface topography as a nonstationary random process*, Nature, Vol.271, 2 February, 1978, pp.431-434.

[21] D.J. Mulvaney, D.E. Newland, K.F. Gill, *A Complete Description of Surface Texture Profiles*, Wear 132 (1989) pp.173-182.

[22] J.A. Ogilvy, *Theory of Wave Scattering from Random Rough Surfaces*, Publisher Adam Hilger, 1991.

[23] P. Kube, A.P. Pentland, *On the Imaging of Fractal Surfaces*, IEEE Trans. Pattern Analysis and Machine Intelligence, 10(5) pp.704-707, 1988.

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

[25] M.L. Smith, *The analysis of surface texture using photometric stereo acquisition and gradient space domain mapping.* Image and Vision Computing, Vol.17 (1999) pp.1009-1019.

[26] K.V. Rajaram, G. Parthasarth, M.A. Faruqui, *A Neural Network Approach to Photometric Stereo Inversion of Real-World Reflectance Maps for Extracting 3-D Shapes of Objects*, IEEE Trans. on Systems, Man and Cybernetics, Vol.25, No.9, Sept. 1995, pp.1289-1300.

[27] G. Kay, 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.

[28] K. Fukunaga *Introduction to Statistical Pattern Recognition*, Second Edition. Academic Press Inc.

Classifier	*Point*	*IRIS*	*SiRIS*
User Time (s)	8	151	158

TABLE I

Comparison of CPU times for feature extraction (includes surface estimation.)

Surface	P_{rms}	*rms difference*
Fracture1	0.052	0.007
Fracture2	0.037	0.002
Fracture3	0.037	0.003
Fracture4	0.022	0.002
Lentils	0.026	0.007
Barley	0.041	0.013
Peas	0.034	0.007
Mixture	0.039	0.013
Milled 125μm	0.016	0.031
Milled 63μm	0.012	0.019
Milled 32μm	0.006	0.007
Milled 16μm	0.003	0.004

TABLE II

Consistency of photometric estimation.

Montage	*Point*	*IRIS*	*SiRIS*
Fracture	31.72	24.57	18.31
Deposit	30.87	18.96	12.15
Ground	40.42	21.23	14.20
Ripple	68.83	14.44	12.95

TABLE III

Experiment 1: Misclassification (%).

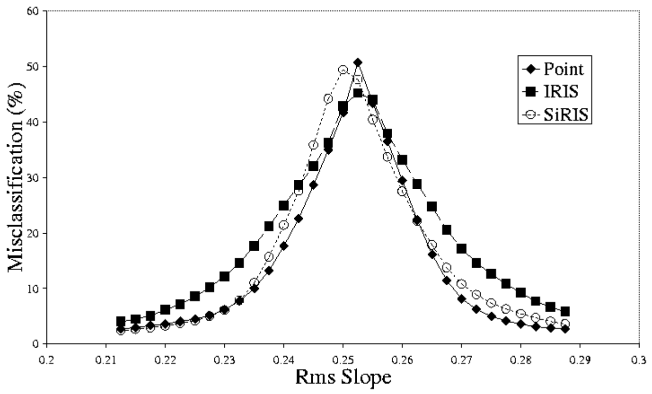

Fig. 8. Simulation 2: Accuracy of classifiers with scaled *Mulvaney* surfaces

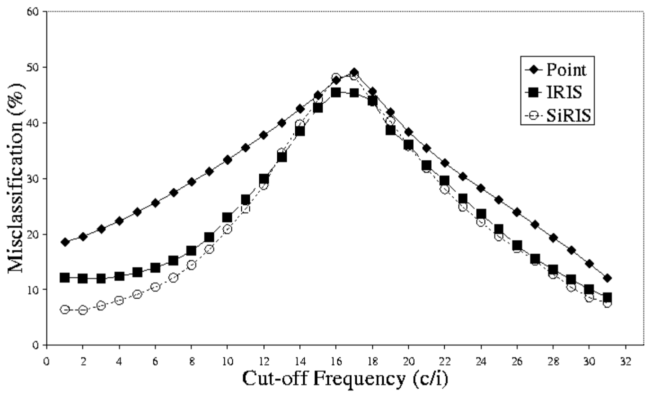

Fig. 9. Simulation 3: Variation of accuracy of classifiers with surface cut-off frequency for Mulvaney surfaces.

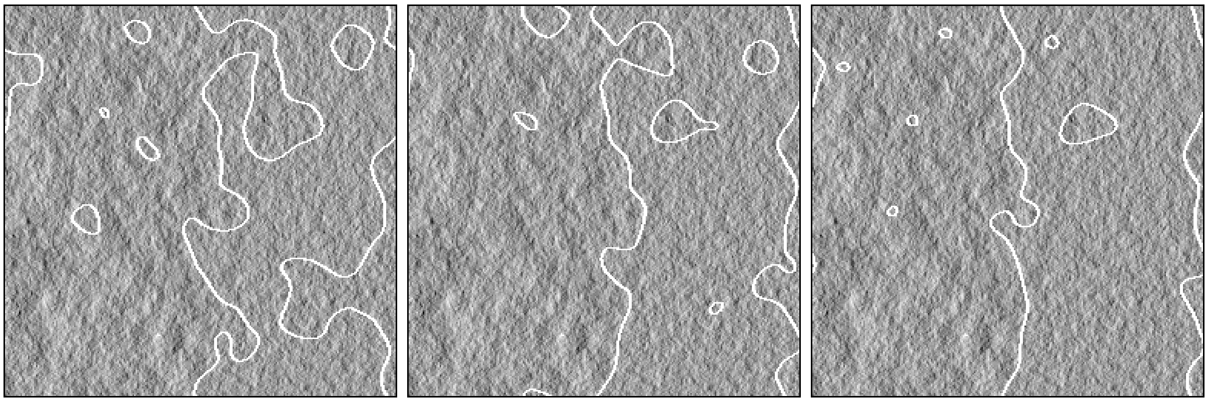

Fig. 10. Simulation 3: Segmentation of *Mulvaney* montage (with cut-off frequencies of 8 and 16 c/i) by Point, IRIS and SiRIS classifiers.

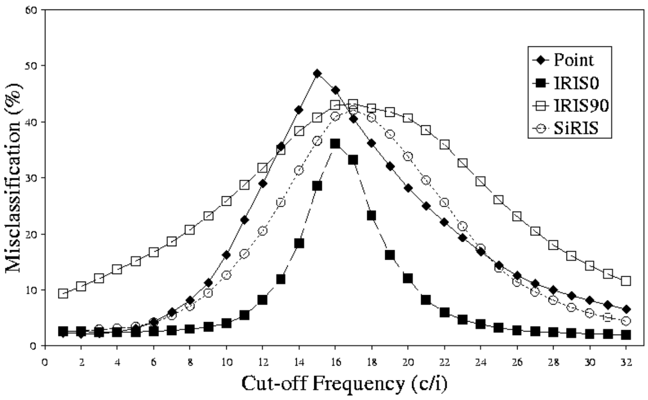

Fig. 11. Simulation 4: Variation of accuracy with horizontal cut-off frequency of *Ogilvy* surfaces.

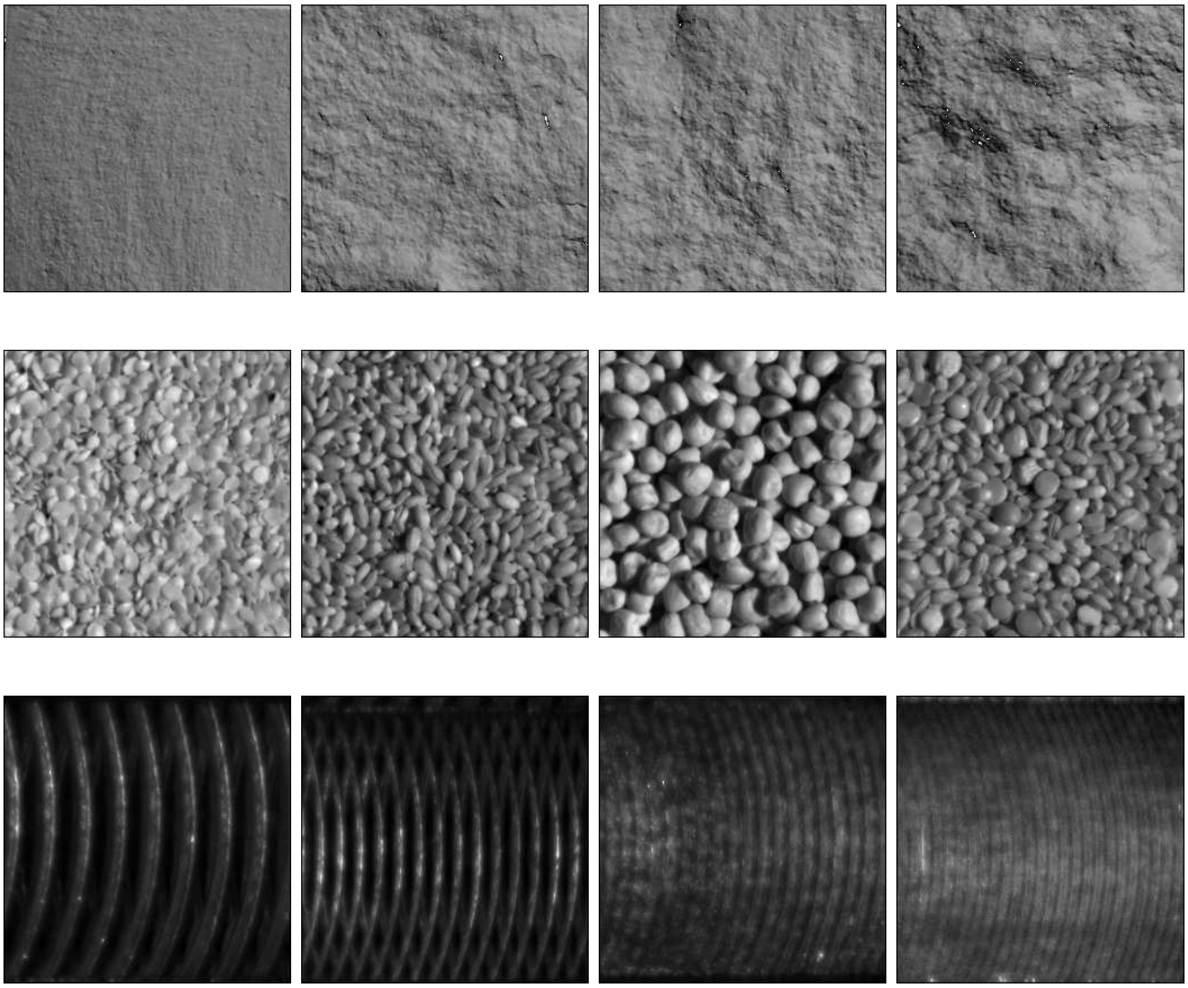

Fig. 12. Test surfaces for consistency experiments: Row 1—Control group; Row 2—Repeating Primitives, Row 3—Milled Surfaces.

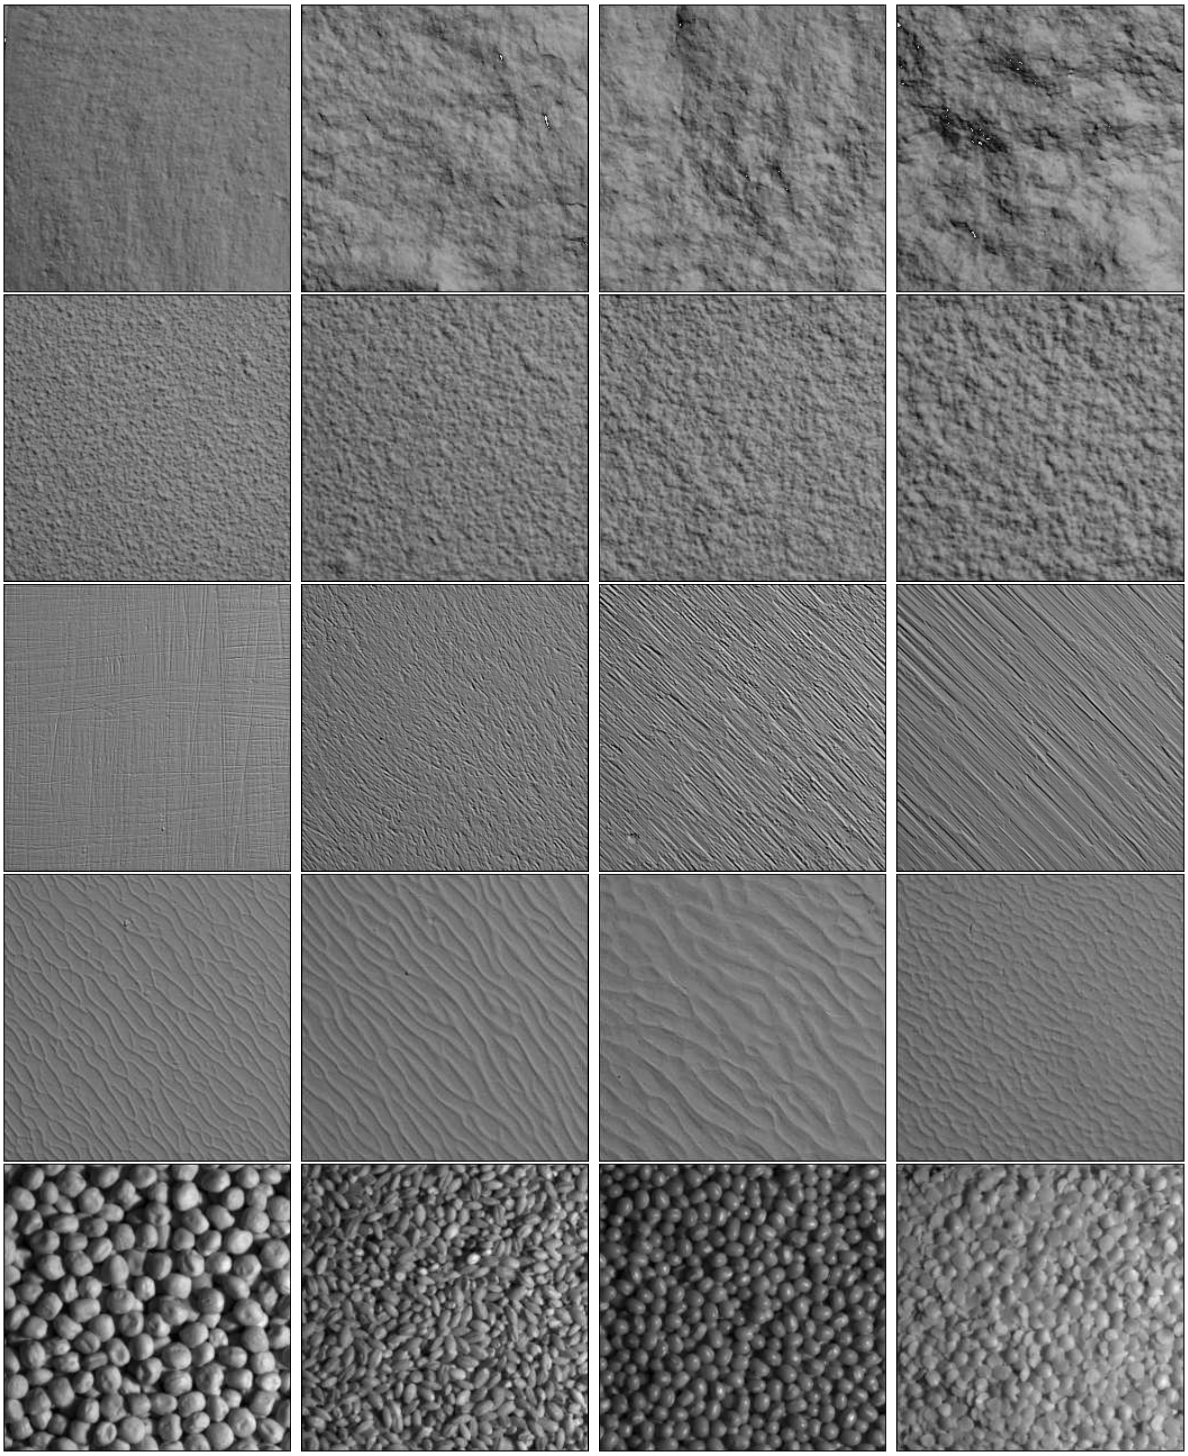

Fig. 13. Classification test surface: Row 1—Fracture surfaces, Row 2—Deposit surfaces, Row 3—Ground Surfaces, Row 4—Ripple surfaces, Row 5—Repeating Primitives.

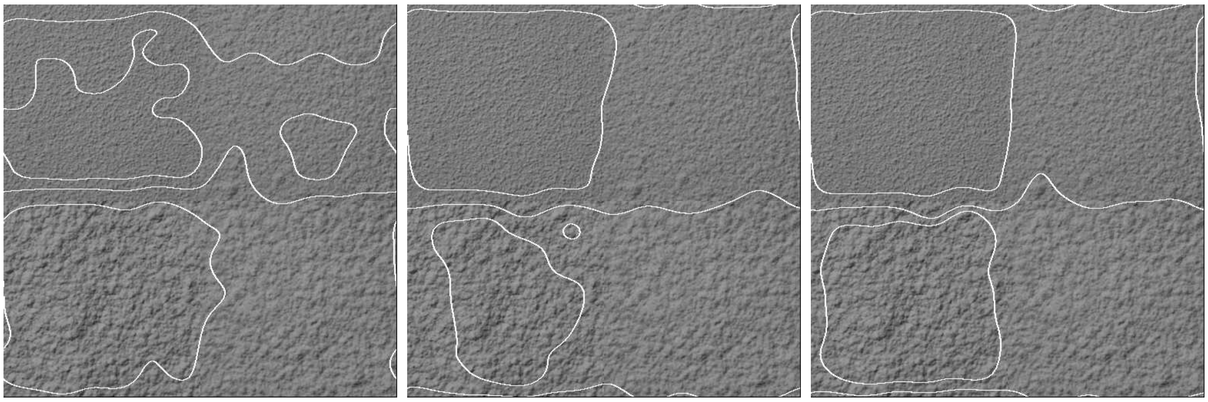

Fig. 14. Experiment 1: Segmentation of *Deposit* montage by Point, IRIS and SiRIS classifiers.

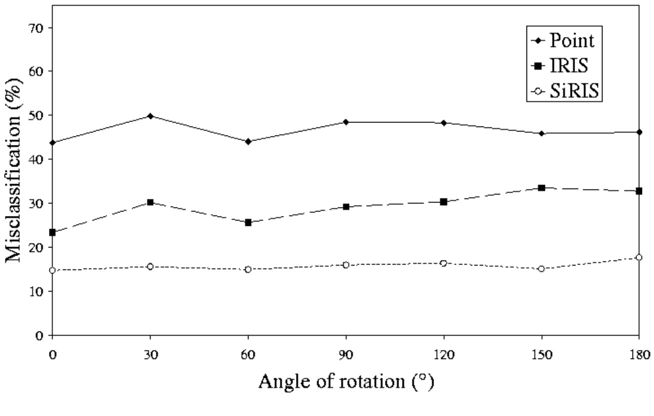

Fig. 15. Experiment 2: Classification accuracy with rotated montage of isotropic surfaces.

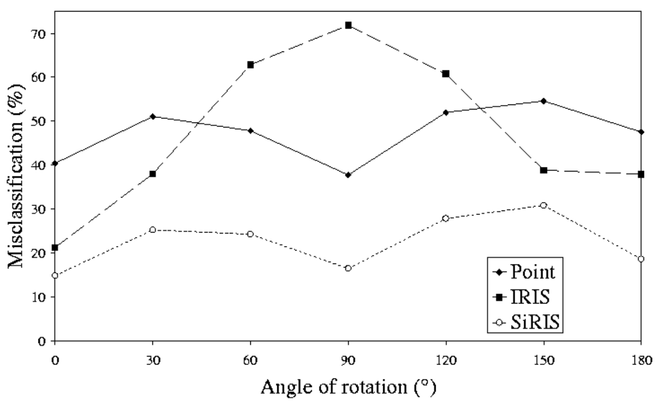

Fig. 16. Experiment 3: Classification accuracy with rotated *Ground* montage.

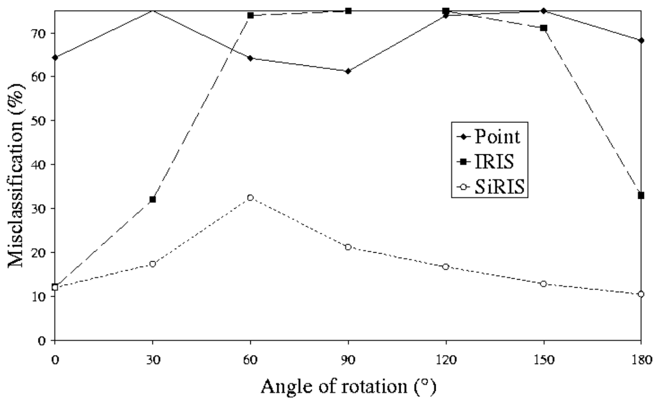

Fig. 17. Experiment 4: Classification accuracy with rotated *Ripple* montage.

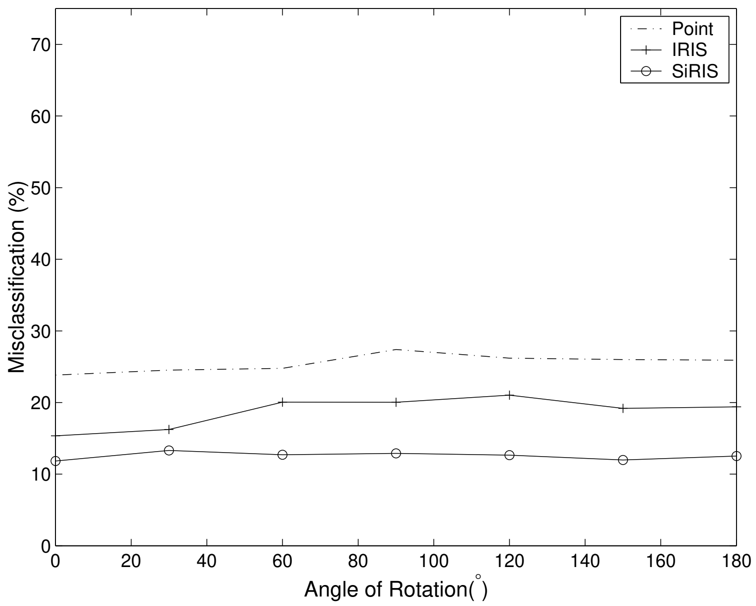

Fig. 18. Experiment 5: Classification accuracy with rotated *Repeating Primitives* montage.