Chapter 9

Summary and Conclusions

9.1 Summary

The aim of this thesis is the development of a classifier that is capable of discriminating between rough surfaces on the basis of their visual appearance, yet which is robust to illuminant tilt. The first part of this aim; that is the development of a rough surface classifier, was undertaken in chapters 2 to 6. A physically-based approach was adopted, which placed the algorithmic classifier in the context of the surface and the intervening processes in the measurement of the data set.

Chapter 2 surveyed methods of surface description adopted in scattering theory and tribology. The power spectrum, in association with a phase condition, was adopted as the principal means of surface description. Three models, framed in terms of this description were adopted. These form the starting point of the simulation work performed in this thesis.

The aim of this thesis requires the classification of rough surfaces on the basis of their visual appearance, and the formation of the image was considered in Chapter 3. Image formation is considered in two distinct stages: the local intensity of a surface facet; and the interaction of facets to form a global image. The local interaction is a function of the surface and illuminant properties, but also of their relative geometeries. Three models of diffuse reflection were compared with the reflectance function for the test surfaces and the classical Lambertian model was adopted.

The global interaction of surface with the illuminant was first modelled by Kube [Kube88]. In the second part of Chapter 3 the accuracy and scope of application of a linear model relating surface and image spectra was considered. The parameter behaviour of the optimal least squares linear filter was compared with that predicted by Kube’s model. The form of Kube’s predictions was found to be accurate for isotropic surfaces and directional surfaces of moderate slope. However, it was noted that there exists an additional scaling relationship, which is related to the degree of surface roughness.

The algorithms adopted in the later chapters of this thesis are not applied to the incident image, rather they are applied to a data set that is subject to noise and distortion. Chapter 4 quantified the nature and magnitude of the imaging process that can be accommodated into the existing model.

Chapter 5 describes the final stage of the surface classifier: the application of the classification algorithms to the data set. This chapter begins by introducing the feature measures and discriminants. An analytical model incorporating feature extraction into the existing model from the previous chapters is developed in Chapter 6. This model predicts that the features will be a function of the illuminant tilt, which implies that the classification will also be affected. The effect of tilt on feature measures, as well as classification was then verified experimentally.

In Chapter 6 the tilt dependency of the classifier was shown; in Chapter 7 several strategies for removing, or at least reducing, this dependency were considered. Three proposals made by Chantler were investigated. The first proposal attempts to cope with tilt induced feature variation by training the classifier over a range of tilt conditions. It was shown that tilt induced variation is much larger than the variation between features of different textures at any given tilt for the test set. This is therefore not an appropriate strategy for our test set. Chantler’s second proposal also trains the classifier under several tilt conditions, however, in this case a series of discriminants are developed which are indexed by tilt angle. The classifier therefore employs whichever discriminant function is closest to the illuminant tilt. While this technique was found to reduce misclassification at the training angles, the misclassification rate quickly rose as the illuminant moved away from this angle. To maintain a reasonable level of classification, up to nine training images would be required. Chantler’s third and favoured approach was based on the inversion of Kube’s frequency domain model. Experimental work in this chapter and in Chapter 3 suggests that the optimal form of the filter varies widely from surface to surface, and the general filter, while reducing the tilt effect to a degree, does not do so to a satisfactory extent.

Since this thesis is concerned with the classification of surfaces from their visual appearance, there would appear to be a great deal of overlap with the field of shape from shading. The central problem of single image shape from shading is that of underdetermination. The majority of algorithms counter this by imposing a smoothness constraint. This is clearly not compatible with our interest in rough surfaces. We have, however, uncovered two techniques in the literature that are more suited to rough surfaces. Pentland's frequency-based scheme [Pentland90] is similar to Chantler's inverse filter, though it recovers a height map rather than a map of the derivative magnitudes. However, this similarity means that the technique will suffer from the same problems as Chantler's technique. Knill uses a pair of linear filters to recover the surface derivatives [Knill90]. Unlike either Pentland's or Chantler's technique, he uses an adaptive algorithm to train the filters. In doing so, he requires *a priori* knowledge of the surface the technique is applied to. By definition this is not available in a classification task.

Our final operation in Chapter 7 was to propose a model-based technique for the suppression of tilt-effects. This used photometric techniques to obtain a model of the surface derivative fields for each training sample. The model was synthetically rendered under the illuminant conditions in which the classifier was to be applied in order to generate appropriate training data.

In Chapter 8 the model-based system was evaluated. This was performed in two stages. The first stage used the models developed in the previous chapters to implement a simulation-based investigation into how well the system recovered rough surface derivatives and predicted the image in the presence of imaging artefacts. The second stage dealt with the evaluation of the system on real data. Although we do not have the means to estimate the accuracy of recovery of rough, real surfaces, it was possible to assess the accuracy for a smooth real surface. It is also possible to measure the accuracy of image and feature prediction. Finally, the performance of a classifier based on the model-based technique was assessed against that of the naive classifier proposed in Chapter 5 and the 'best case' classifier which was retrained under each tilt condition. The model-based technique was found to suppress the effect of tilt variation on classification and in several cases had a mis-classification rate approaching the 'best-case' lower bound.

9.2 Future Work

Slant Invariance

This thesis has dealt only with the problem of tilt invariance and has neglected the effect on classification of variations in slant angle. In theory, a model based scheme should be equally adept at dealing with variations in slant. However, this represents a much more demanding task for the technique since image prediction now requires extrapolation from, rather than interpolation between, the training images. The task is complicated by the fact that the mean level of the texture image will vary with tilt, making camera non-linearity a much more significant problem.

Surface Recovery

In the first chapter of this thesis we proposed a hypothetical inspection task, which requires rough surfaces to be identified. It is not difficult to imagine applications where, having identified an anomaly or other region of interest, it would be useful for a human operator to be able to examine the three dimensional structure of the anomaly. In Chapter 6, we concluded that the single image, rough surface recovery techniques required *a priori* knowledge of the surface type in order to perform well. The identification of the region of interest using classification techniques would allow application of this knowledge to the recovery of surface topography.

Photometric Classification

This thesis has considered the case where the illuminant tilt was not under the control of operator during the classification stage. However, in many cases it will be possible to control the direction of illuminant. Where this is the case, we advocate the integration of photometric techniques into the classification process.

A hypothetical system could use photometric techniques to estimate the surface derivative fields and apply texture analysis techniques to the estimated fields. If the derivative fields could be consistently estimated this would remove the effects of illuminant direction. Furthermore, it would allow the use of much more surface information than is present in the image since the directional filtering effect of the illuminant would be avoided. The removal of the directional filtering would also allow the application of rotation invariant techniques to rough surface textures.

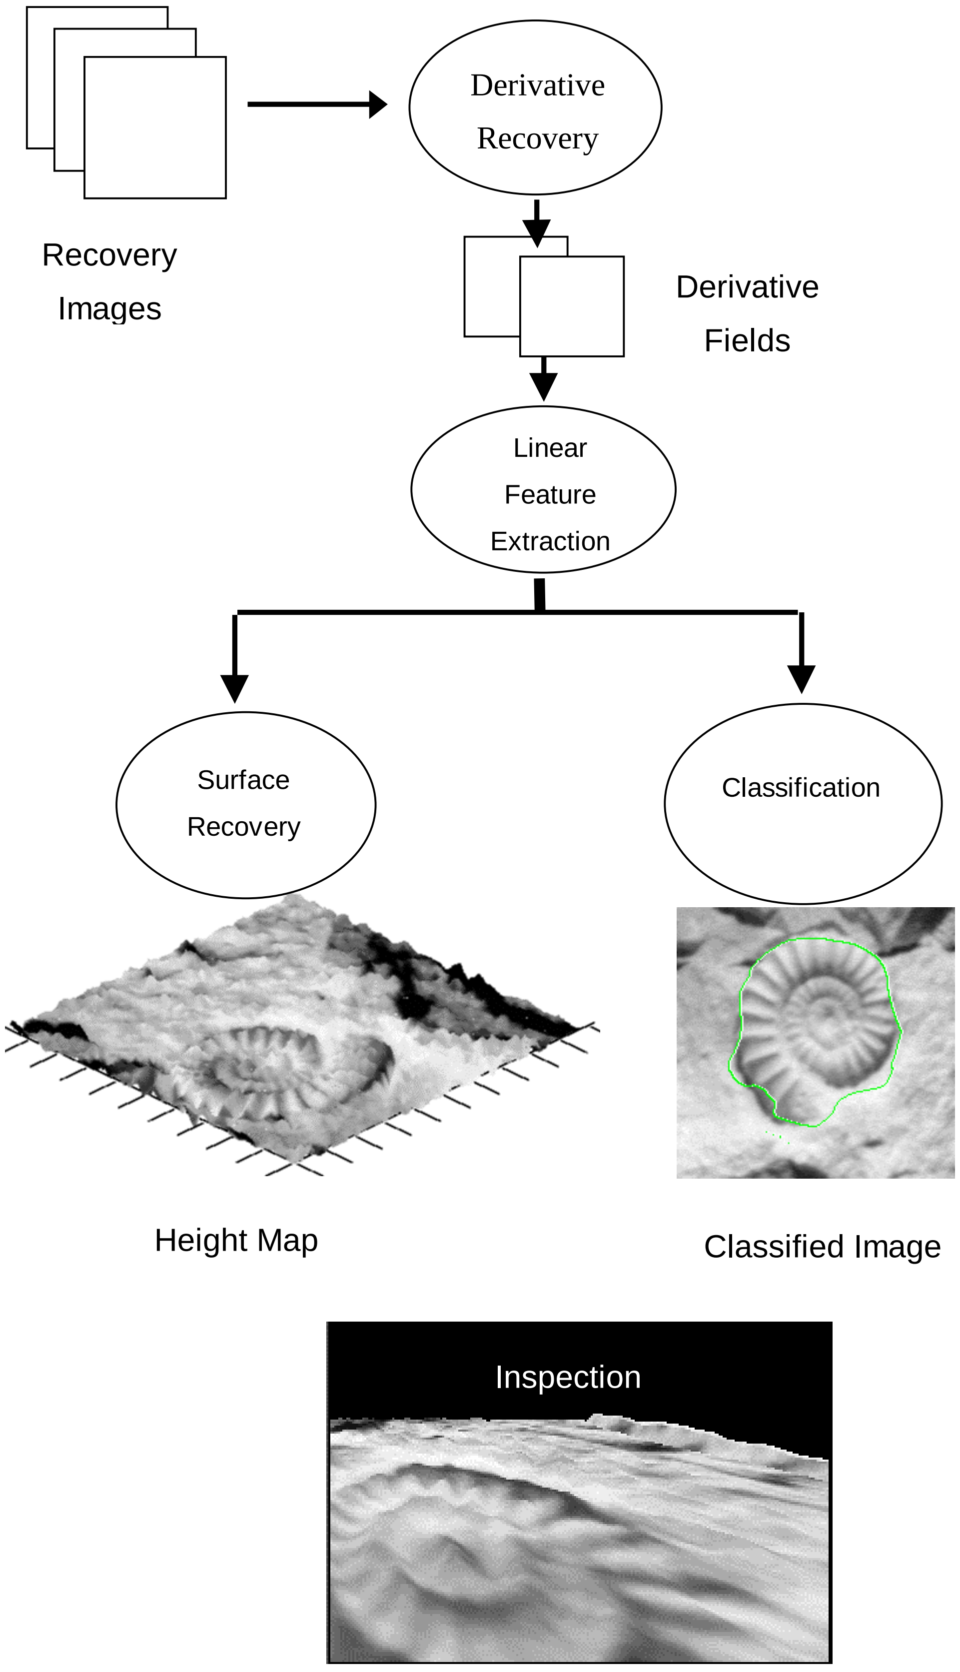

Figure 9.2.1 A proposed scheme integrating classification with surface recovery.

The recovery of derivative fields also suggests that surface recovery is a possibility. This would require the integration of the derivative fields. One significant obstacle is lack of integrability, although this can be overcome relatively easily for smooth surfaces, it is likely to present more serious difficulties for the reconstruction of rough surfaces. While integration can be carried out in the spatial domain, it can also be undertaken in the frequency domain [Frankot88], by weighting the spectral components and swapping the real and imaginary components. Pentland noted that filters localised in space and frequency form an alternative to the Fourier Transform in his frequency based shape from shading algorithm [Pentland90]. Integration of this approach into a photometric classifier would allow the recovery of surface height with little computational cost beyond that already incurred for the classification while avoiding the restrictions on albedo, linearity and noise associated with a single image scheme. The surface type classification, and the use of spectrally defined measures would also allow domain knowledge to be used by an iterative technique. A hypothetical system is shown in *Figure 9.2.1*.

9.3 Conclusions

This thesis has developed a series of models of the process of classifying rough surfaces on the basis of their visual appearance. These models are used in the description, analysis and simulation of the process.

Three spectral models of rough surfaces were adopted from the literature, and their interaction with incident light to form an image was shown in a spectrally definable form. It was found that under certain conditions imaging artefacts could be considered as consisting of a Gaussian blur transfer function, and additive white noise. The combination of the imaging model with the spectral definition of the features showed that the features on which classification is based are functions, not only of the surface, but also of the illuminant tilt angle. This was confirmed by simulation and experiment. The ability of tilt variation to induce classifier failure was demonstrated with both simulation and three montages of real textures.

Several schemes proposed in the literature for reducing the tilt induced effect were evaluated. We concluded that none of these was sufficiently effective for our data set and proposed a new model based scheme. This technique used photometric techniques to estimate the derivative fields of the training surface which could then be rendered under the appropriate illumination conditions to produce training data suited to the prevailing conditions. Evaluation of this scheme showed that it was able to accurately model real textures and to significantly reduce the effect of tilt variation on classification—easily outperforming the other schemes considered.