Chapter 1

Introduction

1.1 Motivation

This thesis is concerned with the application of texture analysis to discriminate between rough, textured, surfaces. In *Figure 1.1.1* we give an example where this ability is required. A fossilised ammonite is embedded in a rock matrix. Both the ammonite and the underlying rock surface have their own, distinct, surface structure which are apparent in the image. This thesis treats the problem of identifying the fossil and isolating it from the background surface as one of detecting a desired signal buried in noise.

Figure 1.1.1 An ammonite (190MYA) fossilised in rock matrix.

Implicit in this approach is the assumption that the observed image of a surface is a function of the surface and is either independent of other variables, or that these variables are constant. In practice, perceived texture may be due to a projected illumination pattern, surface markings, the interaction of illumination with a rough surface or any combination of these. This work is concerned with the automated classification of textures belonging to the third group.

While this type of texture is formed by the interaction of illumination and rough surface, it is desirable to classify the surface type regardless of illumination conditions.

Consequently, work has been carried out on texture classification that is invariant to changes in the illuminant intensity [Thau90] and colour [Healey95]. However, with the exception of [Chantler94], the effect due to variations in illuminant direction is largely neglected in the published literature.

The direction of the illuminant with respect to the texture may be defined by two polar co-ordinates: slant and tilt. Slant is the angle between the camera axis and the illuminant vector; in this work it shall be held constant at 60°. Tilt refers to the polar angle of the illuminant on a plane normal to the camera axis, *Figure 1.1.2a.* The tilt angle (τ) is also shown in *Figure 1.1.2b* with the coplanar general orientation parameter θ. The aim of this work is to maintain consistently accurate classification of textures under changes in the illuminant tilt angle.

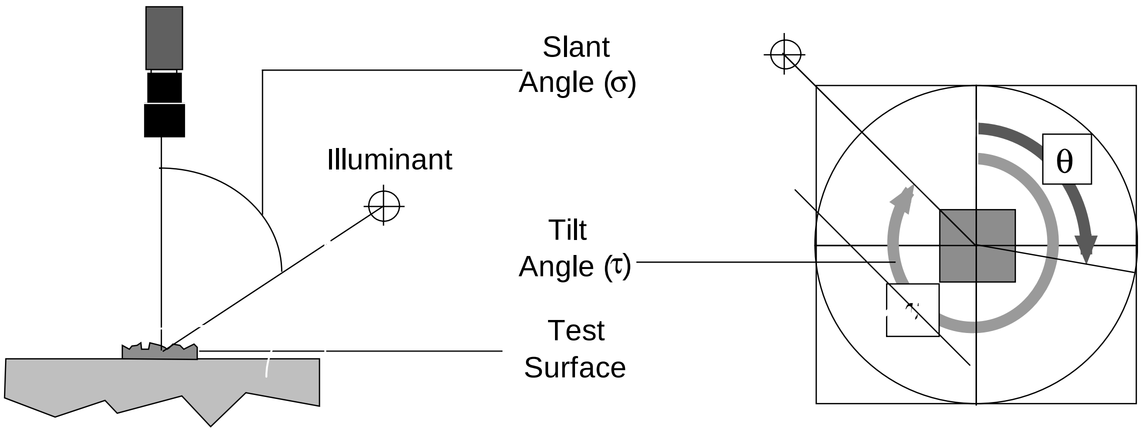

(a.) *(b.)*

Figure 1.1.2 Illuminant tilt and slant angles.

We illustrate the importance of illuminant direction for rough surface discrimination with the following example. Consider the fossilised trilobite (*Elrathia Kingii, 550 MYA, Utah*) shown in *Figure 1.1.3*a. Here the fossil is illuminated from tau=90° and there is little texture information that can be used to segment the fossil from the surrounding matrix. Illuminated from 0° (*Figure 1.1.3b*), however, the animal's pleurae become visible as vertical undulations, allowing the fossil's thorax to be identified using a vertically oriented Gabor filter (*Figure 1.1.3c).*

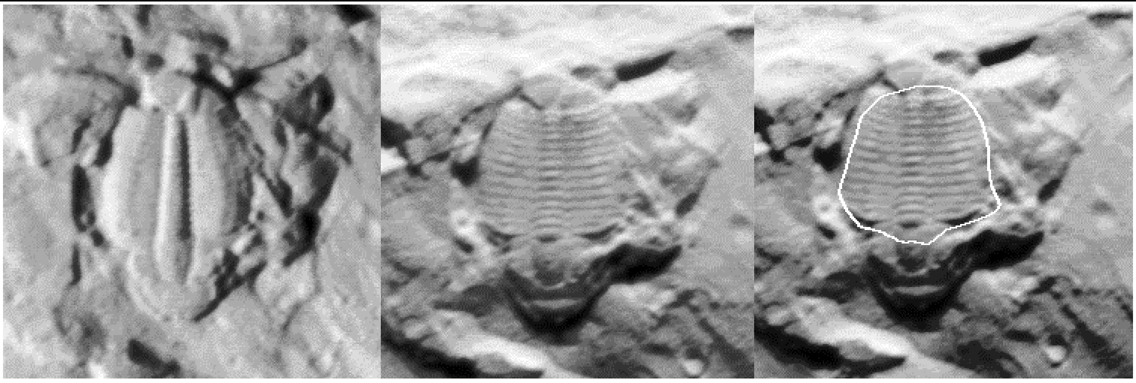

(a)Fossil Illuminated from Tau 90°.	(b)Fossil illuminated from Tau 0°.	(c) The segmented image.

Figure 1.1.3 Classification of a trilobite fossil (550 MYA).

Illuminant tilt is evidently important in attenuating or accentuating directional information that may be critical to classification. As a consequence of this, a shift in the tilt direction will alter the image texture of a given surface. It follows that a classifier which has been trained using images of surfaces obtained at a given tilt angle—which classifies well under these conditions—may not be able to accurately classify the same surfaces illuminated from different tilt angles. The goal of this thesis is to develop a classifier which will be able to classify a surface regardless of the illuminant tilt angle at which the surface was imaged.

1.2 Texture Analysis: A Brief Overview

In this thesis we will use techniques developed in texture analysis to discriminate between rough surfaces based on their visual appearance. Texture analysis, itself, has a long lineage and represents the confluence of several disciplines. In this section we briefly note the main areas which have provided either the motivation to investigate texture, or the techniques which have been used to do so.

Biological visual systems are biased towards the detection of rapid changes of intensity over either time or space, since these stimuli are most likely to be important to the organism. As a consequence of the sensitivity to spatial variation, textures are particularly useful stimuli for the investigation of visual systems. Furthermore, as Gibson stressed, biological visual systems are adept at inferring surface orientation from texture cues [Gibson50]. In fact, Gibson saw texture as being more than a visual cue but rather as having a fundamental role in his ecological view of perception[Bruce p.226]. As Bruce and Green put it "*the input for a receptor is a stream of photons, but the input for a perceiver is a pattern of light extending over space and time*" [Bruce p.376]. Gibson's work analyses the relationship between the optical pattern and the observer's environment. Texture and optical flow are considered the primary sources of information about the environment [Gibson79]. Perhaps the most convincing evidence of the biological importance of texture is its ability to affect visual systems using disruptive pattern camouflage, where texture is used to 'jam' or drown out visual cues such as outlines.

Biological systems have also contributed to texture analysis in a less predictable way. The use of a disruptive camouflage pattern requires pigmentation cells to 'switch' to one of two or more colourings. The cells are identical, and have no centralised form of communication, i.e. they have no sense of their position, yet they must switch colouring in a spatially organised manner. Turing hypothesized that the mechanism by which they do so is due to the concentration of levels of certain chemicals. The processes by which the chemicals diffuse across the surface and react with each other give rise to the global spatial pattern [Turing52]. Recently, several algorithms that simulate the reaction diffusion effect to produce textures have been developed for computer graphics [Turk90][Witkin90]. Picard includes reaction diffusion as an important class of textures in her, "society of (texture) models", and notes its ability to synthesize textures comprised of spots and stripes [Picard96].

In a typically sized image there may be approximately a quarter of a million data elements—an equivalent volume of data to a census of a small city; consequently it is hardly surprising that statistical methods should have been applied. What is special about texture images, from the statistical point of view, is the spatial interaction of both neighbouring and distant elements. Texture images have allowed the investigation of spatial interaction models in a much more visual and intuitive way. The spatial interaction may be described in terms of second order probabilities and several methods have been developed which attempt to parsimoniously extract the information held in the joint distributions. Co-occurence techniques e.g.[Haralick73] extract features from the distributions, whereas Markov models e.g.[Chellappa85] have been used to model the second order probabilities. In either case the algorithms are generally limited to observation of the interactions between immmediately neighbouring pixels.

A large class of textured images may be considered as two-dimensional stochastic signals and consequently fall into the realm of random signal processing. Analogously to the statistical approach, the interaction and interdependence of neighbouring pixels has motivated the use of techniques which are capable of predicting the likely value of a pixel's intensity from those of its neighbours. Autoregressive models have been used, both on their own and with Moving Average models, to both model and discriminate textures [Kayshap84]. Furthermore, the requirement to localise what is an essentially non-local phenomenon has motivated the use of time/frequency analysis. One of the main areas of research has been the generalisation of one dimensional signal processing algorithms to two dimensions. Wavelets [Livens97], Wigner-Ville distributions [Song92] and many other techniques have followed this path.

With the partial exception of those techniques that use texture as a cue to orientation, all the approaches mentioned above take the data set as their starting point. In this thesis we shall apply texture analysis techniques to a data set which does have a specified physical meaning—that of a rough surface interacting with light. By treating the imaged texture as a function of surface topography, which will be used to discriminate between surfaces, we must account for certain variables in the physical process that may affect the data set and our classification of it. The goal of this thesis is to develop a classifier that is immune to variation of one of these variables: *illuminant tilt.*

1.3 Scope

In defining the scope of this research it is useful to adopt a hypothetical example: consider an inspection system tasked with segmenting a rough, textured, surface into regions with uniform surface characteristics. This will be accomplished using an overhead CCD camera, directional lighting and an algorithmic classification system (*Figure 1.3.1*).

The surface is considered to be globally flat, i.e. each of its partial derivative fields sum approximately to zero. Furthermore, the surface reflectance function is considered to be diffuse, and uniform throughout the surface. The light rays incident on the surface are considered to be parallel and of equal intensity across the imaged region. This is consistent with a point source located at infinity or light originating from a parabolic reflector. The imaging device is assumed to be a monochrome CCD camera located directly overhead the texture sample. In addition, we shall assume the topography of the surface is small relative to the distance between the camera and the surface, and that the projection of the surface onto the CCD array is orthographic.

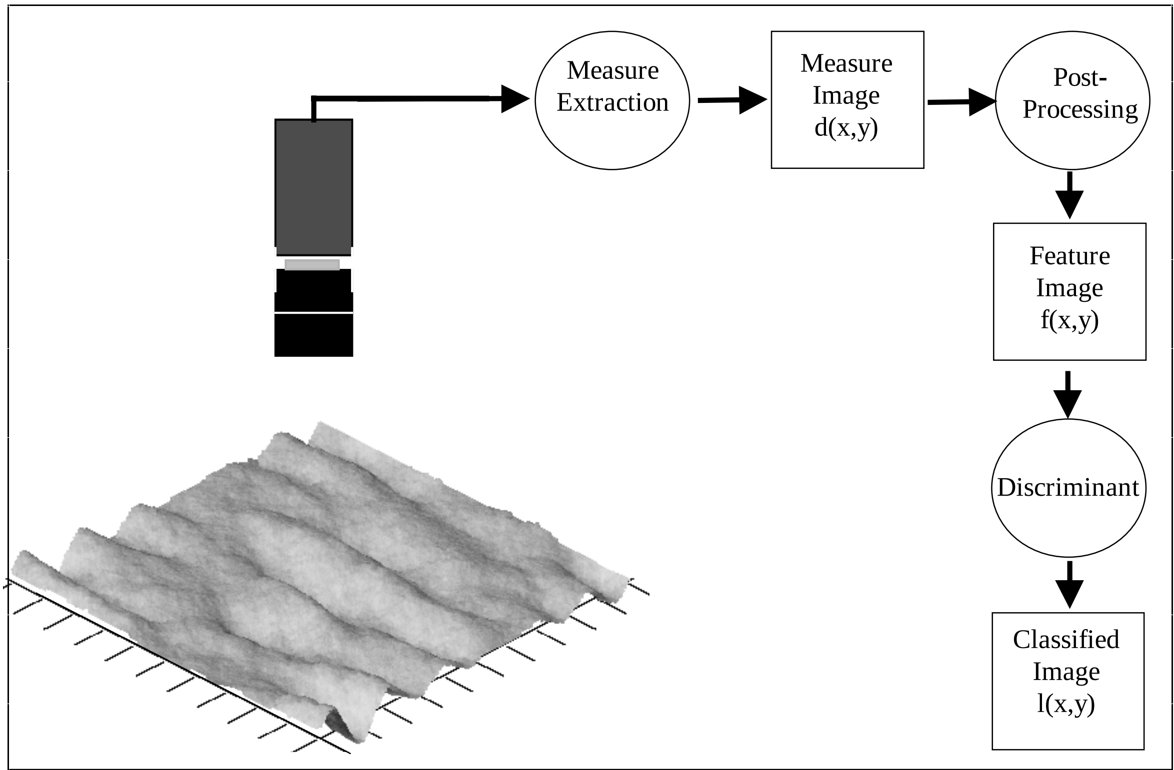

Figure 1.3.1 The classification system.

Whereas most work in texture analysis has taken the stored data set as its starting point, the work described here is concerned with modelling the transitions from surface to image to numerical and finally symbolic representation. This thesis is specifically concerned with modelling and remedying the effect of changes in the direction of the illuminant tilt for the purposes of surface classification.

1.4 Original Work

This thesis makes three main contributions to the field of texture analysis:

- The first contribution of this thesis is the combination of physically verified models, drawn from several areas within the field of computer vision and others external to it. The combination of these elements allows a much more analytical approach to texture analysis within the scope of this thesis. The analytical approach advocated in this thesis allows the systematic development and evaluation of texture analysis algorithms in the context of the actual physical system.The application of physics-based vision techniques was pioneered in the field of shape from shading, however, its application to texture analysis, as described here, is original. Pentland published a series of papers linking fractal surfaces to observed textures

[Pentland84,85,86 and Kube88]. He did not present any physical justification for the models of the surface or the reflectance function he used. Healey has been active in many of the fields considered in this thesis, such as reflectance [Healey89], imaging artefacts [Healey94] and texture analysis [Speis96]; he has not, however, attempted to relate these areas to each other.

- We have described the effect of illuminant tilt on a set of Gabor features and shown that this effect may introduce a significant degree of misclassification. The tilt induced misclassification problem was first noted by Chantler [Chantler94]. By adopting feature measures that are consistent with the surface and image models, we believe we have both verified many of Chantler's findings and also formed expressions that are more coherent and intuitively understandable.
- An important contribution of this thesis is the development of a model-based technique to reduce the effect of tilt changes on classifier accuracy. We have verified the performance of the system for tilt variation and believe it offers an effective and physically meaningful mechanism for the suppression of tilt effects.

The combined result of these contributions may be stated as follows; a physical phenomenon was accurately predicted and the effects at the algorithmic level modelled. Using our understanding of the physical system, a technique was developed which reduces the effect of the physical phenomena on the computional level. In this way, the goal of the thesis, i.e. to develop a classifier robust to illuminant tilt, is achieved.

1.5 Thesis Organisation

This thesis may be divided into two chapter groupings. The first group comprises Chapters 2 to 6 and develops a theoretical framework for the classification of rough, textured, surfaces on the basis of their visual appearance. One of the predictions of the model is that the visual texture, as well as being a function of the surface, will be affected by the illuminant tilt. The second part of this thesis, Chapters 7 and 8, focuses on techniques for the consistent classification of rough surfaces robust to illuminant direction.

Development of an analytical model is broken into the chapter aims and objectives shown in *Table 1.5.1*. Analysis begins in Chapter 2 with the rough surface. The literature associated with tribology and scattering theory is consulted to obtain firstly, a means of describing rough surfaces, and secondly, models of rough surfaces framed in terms of that form of description. In Chapter 3 a spectral model of the relationship between the surface description and the image incident on the camera lens is developed. An existing model [Kube88] is evaluated in terms of optimality and of accuracy. The conversion of this image into a usable data set is carried out by camera and framestore, and is modelled in Chapter 4. This chapter forms an analytical model of the imaged data and develops a noise engine with estimated parameters of the conversion process. The classifier developed in Chapter 5 operates on the data set and converts the two dimensional signal into symbolic form. This represents the final component in the conversion from physical surface to symbolic representation. However, in Chapter 6 it is shown that varying illuminant direction between training and classification can render the classifier useless; the second part of this thesis examines techniques designed to remove this dependency.

Thesis Goal	*To develop a tilt invariant classifier*				
Thesis Sub-goal	*To form a model of the effect of tilt*				
Chapter	Chapter 2	Chapter 3	Chapter 4	Chapter 5	Chapter 6
Aims	To model the topography of rough surfaces.	To model the process of image formation.	To model the imaging device.	To develop a texture classifier suited to rough surfaces.	To model the classifier.
Objectives	To find: • a means of description • models of rough surfaces	• To describe the local effect • To develop a spectral model.	To model: • The device transfer function • Sources of noise	Selection of appropriate algorithms for: • Measure extraction • Post-Processing • Discrimination	Modelling the effect of tilt on: • Measure spectra • Feature statistics

Table 1.5.1 Aims and objectives of first chapter grouping.

The second chapter grouping considers several schemes designed to tackle the problem of tilt dependency. Chapter 7 considers three techniques advanced by Chantler [Chantler94] as well as the feasibility of applying shape from shading techniques. In fact, whereas some of these schemes do reduce the effect of tilt on classification, none eliminates the effect. In Chapter 8 a novel technique is proposed, which models the surface and the processes analysed in the first part of this thesis during its training stage. The technique is evaluated by simulation using various model surfaces, the empirical reflectance map and the noise engine. The scheme's performance on real data is then evaluated at each process step before final evaluation on the criterion of misclassification.

Finally, Chapter 9 summarises the findings of the previous chapters and draws together the conclusions of the work carried out in this thesis.