

**Surface Reflection Model Estimation
from Naturally Illuminated Image
Sequences**

by

Robert Charles Love

*School of Computer Studies
The University of Leeds*

September 1997



Submitted in accordance with the requirements
for the degree of Doctor of Philosophy.

The candidate confirms that the work submitted is his own and that appropriate credit
has been given where reference has been made to the work of others.

Surface Reflection Model Estimation from Naturally Illuminated Image Sequences

Abstract

This thesis addresses the problem of estimating the surface reflection model of objects observed in a terrestrial scene, illuminated by natural illumination; that is, a scene which is illuminated by sun and sky light alone. This is a departure from the traditional analysis of laboratory scenes, which are illuminated by idealised light sources with positions and radiance distributions that are precisely controlled. Natural illumination presents a complex hemispherical light source which changes in both spatial and spectral distribution with time, terrestrial location, and atmospheric conditions.

An image-based approach to the measurement of surface reflection is presented. The use of a sequence of images, taken over a period of time, allows the varying reflection from the scene due to the changing natural illumination to be measured. It is shown that the temporal change in image pixel values is suitable for the parameters of a reflection model to be estimated. These parameters are estimated using regression techniques. Two such regression methods are considered: a traditional non-linear method and the probabilistic approach of simulated annealing. It is shown that simulated annealing provides consistent performance in this application.

This work focuses on the use of physically-based models of illumination, surface reflection and camera response. Using such models allows the system to produce quantitative, as opposed to qualitative, results and allows radiometric measurements to be made from image pixel values. The use of accurate models of daylight illumination allows scenes illuminated by skies of varying atmospheric conditions to be considered. The results obtained by the presented methods may be used for a variety of tasks ranging from object recognition to the automated generation of virtual environments.

Results are presented which show that the proposed method is suitable for the wide variety of camera positions, surface orientations and sky conditions that may be experienced. The method is also shown to be tolerant of image noise and may be used on single or multiple pixels within each image. These results are based on the analysis of synthetic image sequences generated using a validated lighting simulation system. Results are also presented for real data recorded using a camera.

Acknowledgements

I would like to thank the many people who have helped in this endeavour. Thanks are due to the members of the Vision Group, both past and present, in addition to the members of the School who have helped to make my time in Leeds both enjoyable and rewarding.

Special thanks are also due to my family and friends who have provided financial and moral support throughout the highs and lows that this research has presented me.

For their encouragement and supervision during this course of work, and especially for their time spent proof-reading this thesis, I would like to thank Nick Efford and Roger Boyle.

I would also like to thank Greg Ward (Silicon Graphics, Inc) for his assistance with the *Radiance* lighting simulation system.

The financial support of the Engineering and Physical Science Research Council is gratefully acknowledged.

Contents

1	Introduction	1
1.1	Quantifying Reflection	2
1.2	Physics-Based Machine Vision	3
1.3	Approach Used	4
1.4	Overview of the Thesis	6
2	Related Work	7
2.1	Image Based Measurement of BRDF	8
2.2	Greyscale Image Analysis	10
2.2.1	Single Image Methods	10
2.2.2	Multiple Image Methods	13
2.3	Colour Image Analysis	15
2.3.1	Colour Constancy	16
2.3.2	Colour Model Estimation	17
2.4	Summary	21
3	Surface Reflection	22
3.1	Reflectance: The Measure of Reflection	23
3.1.1	Reflectance	24
3.1.2	The Bidirectional Reflectance Distribution Function	26
3.2	Fresnel’s Laws of Reflection	27
3.3	Modelling Surface Structure	28

3.3.1	Surface Height Distribution	29
3.3.2	Facet Slope Distribution	30
3.3.3	Shadowing and Masking	31
3.4	Modelling Surface Reflectance	31
3.4.1	Ideal Specular Reflection	34
3.4.2	The Normal Lobe	34
3.4.3	The Forescatter Lobe	35
3.4.4	The Backscatter Lobe	36
3.4.5	Ward's Reflection Model	36
3.5	Summary	38
4	Natural Illumination	40
4.1	Direct Solar Illumination	41
4.1.1	Solar Geometry	41
4.1.2	Solar Intensity	43
4.1.3	Spectral Distribution	45
4.2	Diffuse Sky Light Illumination	47
4.2.1	Sky Light Distribution Models	48
4.2.2	Sky Light Luminous Efficacy	53
4.2.3	Spectral Distribution	53
4.2.4	Surface Irradiance	54
4.3	Summary	59
5	A Sensor Model	60
5.1	Projection of the Scene	61
5.2	Image Plane Irradiance	61
5.3	Production of the Digital Image	63
5.3.1	Achromatic Image Formation	64
5.3.2	Colour Image Formation	65

5.4	Sources of Error and Calibration	66
5.4.1	Optical Errors	67
5.4.2	Sensor Errors	67
5.5	Summary	68
6	Surface Model Estimation	69
6.1	The χ^2 Metric	70
6.2	Description of Data and Associated Model	72
6.2.1	Selection of a Suitable Model	74
6.2.2	χ^2 for an Image Sequence	75
6.3	Topology of the χ^2 Hypersurface	77
6.3.1	Effect of Camera Position	79
6.3.2	Effect of Sky Conditions	81
6.3.3	Effect of Camera Noise	83
6.3.4	Sampling Multiple Pixels with Similar Surface Normals	85
6.3.5	Sampling Multiple Pixels with Differing Surface Normals	86
6.3.6	Effect of Sequence Length and Image Frequency	89
6.3.7	Observations from Simple Experiments	91
6.4	Minimising χ^2	92
6.4.1	Brute Force Search	92
6.4.2	Gradient Search	93
6.4.3	Analytical Search	93
6.4.4	Simulated Annealing	95
6.4.5	Limiting the Search Space	98
6.5	Summary	99
7	Results	100
7.1	Overview of the Reflection Model	101
7.2	Levenberg-Marquardt Regression	102

7.2.1	Two-Parameter Reflection Model	102
7.2.2	Achromatic Reflection Model	107
7.2.3	Summary of Levenberg-Marquardt Performance	107
7.3	Simulated Annealing	109
7.3.1	Selection of Cooling Schedule	110
7.3.2	Two-Parameter Reflection Model	112
7.3.3	Achromatic Reflection Model	113
7.3.4	Colour Reflection Model	116
7.3.5	Summary of Simulated Annealing Performance	118
7.4	Real Image Sequence	118
7.4.1	Achromatic Analysis	120
7.4.2	Colour Analysis	120
7.5	Summary	123
8	Conclusions	126
8.1	Performance	128
8.2	Future Work	129
8.3	Closing Comments	131
A	Radiometry and Photometry for Machine Vision	132

List of Tables

3.1	The nine reflectance geometries	25
7.1	Two-parameter Levenberg-Marquardt results (Lambertian surface)	103
7.2	Two-parameter Levenberg-Marquardt results (rough metallic surface)	104
7.3	Two-parameter Levenberg-Marquardt results (smooth metallic surface)	106
7.4	Two-parameter simulated annealing results (Lambertian surface)	113
7.5	Two-parameter simulated annealing results (rough metallic surface)	113
7.6	Two-parameter simulated annealing results (smooth metallic surface)	114
7.7	Three-parameter simulated annealing results (smooth plastic surface)	115
7.8	Three-parameter simulated annealing results (rough metallic surface)	115
7.9	Three-parameter simulated annealing results (smooth metallic surface)	116
7.10	Five-parameter simulated annealing results (blue plastic surface)	117
7.11	Five-parameter simulated annealing results (green metallic surface)	117
7.12	Five-parameter simulated annealing results (white metallic surface)	117
7.13	Description of sample materials	119
7.14	Estimated achromatic model parameters for material samples	120
7.15	Estimated colour model parameters for material samples	122

List of Figures

1.1	Example image sequence	5
3.1	Surface reflection geometry	23
3.2	Behaviour of an incident wave at the interface between two media	27
3.3	Difference in path lengths for light reflected from a rough surface	29
3.4	Surface height distribution model	29
3.5	Facet slope distribution model	30
3.6	Facet shadowing and masking	31
3.7	Linear approximation to BRDF	32
3.8	Light reflection from pigmented dielectric material	32
4.1	Solar geometry	42
4.2	Extraterrestrial solar irradiance incident on a surface	44
4.3	Tilted surface geometry	45
4.4	Solar irradiance on a horizontal and tilted surface	46
4.5	Sky luminance geometry	48
4.6	Clear sky luminance relative to zenith	50
4.7	Sky zenith luminance for clear and overcast skies	50
4.8	Overcast sky luminance relative to zenith	51
4.9	Intermediate sky luminance relative to zenith	52
4.10	Irradiance from a CIE clear sky calculated on the surface of a sphere	55
4.11	Use of a local hemisphere to sample the sky hemisphere	55
4.12	Geometry for Hall's hemispherical integrator	56

4.13	Geometry for Nishita's integration method	57
4.14	Reconstructed image sequence	59
5.1	Imaging geometry for a simple camera model	62
5.2	Relative sensitivity to wavelength of a CCD camera and the human eye	65
5.3	Relative sensitivity of CCD camera to filtered wavelengths	66
6.1	Sensor measuring reflected radiance within a small solid angle	72
6.2	Methods for obtaining further measurements of surface reflection	73
6.3	The change in image pixel value across a sequence	74
6.4	Selected images from a sequence showing a Lambertian disc	78
6.5	χ^2 function for a clear sky, Lambertian scene viewed from the north	79
6.6	χ^2 function for a clear sky scene viewed from the north	79
6.7	χ^2 function for a clear sky scene viewed from the east	80
6.8	χ^2 function for a clear sky scene viewed from the south	80
6.9	χ^2 function for an intermediate sky scene	82
6.10	χ^2 function for an overcast sky scene	82
6.11	Value of a pixel in each image of a noisy sequence	83
6.12	χ^2 function for an image sequence with 5% noise	84
6.13	χ^2 function for an image sequence with 10% noise	84
6.14	Values of four pixels within an image across a sequence	85
6.15	χ^2 function for four sampled pixels in an image sequence	86
6.16	Selected images from a sequence showing a Lambertian sphere	87
6.17	Values of five pixels within an image sequence	88
6.18	χ^2 function for sequence observing sphere under a clear sky from the east	88
6.19	χ^2 function for sequence observing sphere under a clear sky from the south	89
6.20	χ^2 function for sequence observing sphere under an overcast sky	89
6.21	χ^2 function for one pixel in a short image sequence	90
6.22	Sampling the temporal change in reflection	91

6.23	Simplex translations	97
6.24	Phase transition region	98
7.1	Two-parameter Levenberg-Marquardt error (Lambertian surface)	104
7.2	Two-parameter Levenberg-Marquardt error (rough metallic surface)	105
7.3	Two-parameter Levenberg-Marquardt error (smooth metallic surface)	106
7.4	Variance of two-parameter estimates	107
7.5	Three-parameter Levenberg-Marquardt error (smooth plastic surface)	108
7.6	Three-parameter Levenberg-Marquardt error (rough metallic surface)	108
7.7	Three-parameter Levenberg-Marquardt error (rough metallic surface)	109
7.8	Effect of initial temperature on simulated annealing performance	111
7.9	Effect of cooling rate on simulated annealing performance	112
7.10	Typical random walk over χ^2 surface	114
7.11	Image selected from real image sequence	119
7.12	Single pixel values for real image sequence	121
7.13	Reconstruction of real image sequence	123
7.14	Modelled pixel values for real image sequence	124
A.1	Spectral luminous efficacy curve	133

Chapter 1

Introduction

A human observing a scene can acquire a wealth of information regarding the shape of objects in the scene, the nature of the light sources and the nature of the material with which objects are composed. This in-depth analysis of a scene, possible even when presented with a single image, is attained with little conscious effort by the observer. To date, no machine-based system has managed to achieve a performance similar to that of the human visual system. This thesis focuses on a single aspect of these abilities: analysis of the surface material or, more specifically, the modelling of light reflection from that material. However, this cannot be achieved without consideration of object shape and illumination. The appearance of an object is a direct product of object shape, surface material and illumination. None of these may be adequately considered, in the context of image shading, without accounting for the accompanying factors.

This thesis focuses on the estimation of a surface reflection model for an object within a naturally illuminated scene. That is, a scene which is illuminated by light from the sun and sky. This represents a dramatic departure from the traditional analysis of images which have been obtained under laboratory conditions using idealised light sources. Such ideal sources of illumination provide convenient conditions in which to interpret image information. Natural illumination, however, provides a complex and ever changing light source. This change is both temporal and spatial. The apparent movement of the sun with time yields an illumination which changes throughout the day, the year, and for differing terrestrial locations. The illumination due to the sky changes not only with the motion of the sun but also with atmospheric conditions. These conditions provide a hemispherical source of illumination which exhibits spatial variation in brightness and spectral

composition.

Results from the analysis of naturally illuminated scenes can be used for a variety of purposes. Object recognition relies heavily on both shape and material properties; using either characteristic in isolation will not provide a generalised recognition system. For example, a robotic system requested to select the “*glossy red ball*” requires both the shape and reflection characteristics of each object in the scene to ensure that the task is performed correctly.

The results may also be used for remote sensing purposes, such as planetary surface exploration. The reflective properties of a planetary terrain provide vital information on the physical structure and chemical composition of its topmost layer, and hence on the geological history of the terrain.

This technique could also be used for the generation of virtual environments based upon existing, naturally illuminated, scenes. Once the reflective properties of the scene have been measured, the resulting virtual environment may be illuminated using a wide variety of illuminants, not just natural ones, and observed from camera positions not used when the scene was analysed. A practical application of this would be in architectural design. The material reflection properties of an existing building may be measured and the resulting information used to model structural modifications or extensions in a virtual environment.

1.1 Quantifying Reflection

The study, and quantification, of light reflection dates back to the 18th century. At this time the performance and limitations of the human visual system were being explored. Leading this pioneering work was Pierre Bouguer (1760). Bouguer’s work defined the field of photometry and hence allowed scientific measurement of human visual capabilities. It is also at this time that the first mathematical models of surface reflection were being formulated. Lambert’s (1760) theory of light scattering from surfaces is still widely used today.

The work of James Clark Maxwell (1864) provided a physical and theoretical framework with which to understand light’s interaction with the environment. The modelling of light as an electromagnetic wave unifies the theory of light propagation with that of radiant heat transfer. This theory, and that of the quantum nature of light, has provided many models of light reflection.

Each of these models presents an approximate solution to Maxwell's fundamental equations, for a variety of surface types.

The field of radiometry provides a physically-based framework within which light reflection may be studied. It provides a set of quantities which allow the reflection from a surface to be measured. This thesis examines two aspects of light reflection: the scattering of incident light by the surface material, and the attenuation of wavelength to produce apparent surface colour. The distribution of light reflected from a surface is determined by the scattering that takes place at the surface and within the bulk of the material. The level of scattering is governed by the material type and the roughness of the surface. The perceived colour of an object is a product of the illumination and the surface material. These factors can be characterised in a functional model of surface reflection.

Using measurements of surface reflection, a mathematical model of a surface's characteristics may be formulated. The resulting model is a function of both illumination and viewing positions as well as the spectral and spatial nature of the light source. A general reflection model may be considered as having variables and parameters. The variables of the model describe the directional quantities being considered and the wavelength of illuminating light. These values will change for a given object as the light source is altered or the object is moved about the scene. However, the parameters of the model characterise the level of scattering from the surface and the wavelength attenuation of incident light. These parameters are fixed for any given surface material, regardless of how the object is viewed or illuminated. It is these parameters which are to be estimated by the system presented in this thesis.

1.2 Physics-Based Machine Vision

In the early 1970s, Berthold Horn saw that the use of radiometric models of light reflection and image formation allowed detailed analysis of digital images. In particular it was shown that surface shape could be estimated from a single image by analysis of image shading. The use of such physically-based models presented a departure from the traditional, geometry-based, approaches which used oversimplified models of illumination, surface reflection and camera optics.

The advent of ever increasing computing power has allowed physically-based models, de-

rived from the consideration of differential quantities, to be employed. These models have been developed in the fields of radiant heat transfer, energy conservation, computer graphics and machine vision itself. Such models are used in this work for all aspects of the imaging process. Whilst accurate models of image formation and surface reflection have been used before in machine vision, the use of physically-based models of daylight illumination have not. This thesis, therefore, presents a significant extension to the study of physics-based machine vision.

The use of such models allows quantitative, as opposed to qualitative, results to be produced. These results may be used in a variety of applications which require accurate measurements. For example, a system which determines a surface to be red and glossy is not providing results suitable for predicting light levels in a room containing the surface. Such calculations require radiometric models of surface reflection that quantify surface properties. This thesis presents a system that can determine appropriate radiometric models for materials in a naturally illuminated scene, given a sequence of images of that scene.

1.3 Approach Used

The approach used in this thesis is to determine surface reflection model parameters for an object observed in a sequence of images obtained with a static camera. The images of the sequence are acquired over a period of time. Given that the observed scene is illuminated by sun and sky light, it is expected that the nature of the illumination will change over the duration of the sequence. Each image of the sequence, therefore, observes the scene in a different set of illumination conditions. A typical sequence of images, taken hourly on a clear day, is shown in Figure 1.1. This sequence clearly shows how the appearance of an object is dependent upon the changing natural illumination in conjunction with the surface properties. It shows not only a change in total brightness but also the spectral change in the illumination.

It is assumed that the location and geometry of the scene, the camera's projection characteristics, and the prevailing weather conditions for each image in the sequence are known *a priori*. Each of these may be estimated using existing methods, which shall be discussed where applicable. Whilst these are significant assumptions, it is felt that future work may look to relaxing these requirements. Further to this, it is shown that the methods developed in this thesis may be used directly if such information is not known, though experiments are not performed to evaluate this.



Figure 1.1: *Sequence of images showing a building, observed hourly, on a clear day.*

The pixel values of an image sequence represent light reflected from the observed surface in the direction of the camera. Using the physically-based models derived for illumination and camera response it is possible to use such pixel values as a measurement of reflected energy. The task presented by this thesis is, therefore, to find the parameters of a model which best characterise the observed reflection. This common problem, of finding the model parameters which best fit observed data, is that of regression. Two methods of regression are considered: Levenberg-Marquardt non-linear regression, and the probabilistic approach of simulated annealing. The former technique is commonly used for such tasks but it is shown to be unsuitable in this case. The use of simulated annealing for this task is a novel approach to solving the problem of reflection model parameter estimation.

This thesis will show that it is possible to estimate the parameters of a surface reflection model which best characterise the object's surface material characteristics using the temporal change in image pixel values. It shall be shown that such analysis may be performed by considering a single or multiple image pixels within each image. The methods suitable for this analysis will be developed by considering synthetic image sequences produced using a validated lighting simulation system. The analysis of such sequences provides an understanding of the limitations of the methods and the effects that the many viewing, illumination, and surface orientation conditions may have on system performance. Using such sequences, the consequences of image noise shall also be considered. Comparative results are presented for the two regression methods considered. Finally, a real image sequence, observing a variety of material samples, and obtained using a camera, is analysed.

1.4 Overview of the Thesis

The remainder of this thesis is structured as follows:

Chapter 2: A review of the relevant background material and related research.

Chapter 3: A detailed discussion of light reflection based upon electromagnetic and quantum theory. A model of surface reflection is presented which accounts for the reflection from a wide variety of surface types.

Chapter 4: A description of daylight illumination models developed by solar energy researchers. This chapter also presents a description and comparison of various methods to enable their use in machine vision tasks.

Chapter 5: Development of a radiometric sensor model which will allow surface reflection to be measured from a colour digital image. The model allows image pixel values to be interpreted in the context of a surface reflecting light which is incident from a possible hemisphere of directions.

Chapter 6: A description of methods which may be used to estimate the parameters of a surface reflection model from a sequence of colour images. These methods are developed by considering a number of synthetic image sequences. These sequences allow the limitations of the methods to be explored for the wide variety of viewing, illumination and surface orientation conditions that may exist.

Chapter 7: Results for two of the suggested regression methods. Results are shown from the analysis of synthetic image sequences in addition to a real image sequence.

Chapter 8: Conclusions and a discussion of future work.

Appendix A: Relevant radiometric and photometric definitions and nomenclature.

Chapter 2

Related Work

The measurement of surface reflection or, more specifically, a bidirectional reflectance-distribution function (BRDF)¹ may be achieved through the use of expensive, specialised equipment such as a gonireflectometer [89]. Such devices use precise light sources, often lasers, and accurately calibrated sensors to measure reflectance for the full range of possible incident and reflection directions. The fine precision of such measurements results in data from which an accurate BRDF model can be developed. Owing to their precise nature and specialised application area these devices are expensive and often inaccessible. There is also the need to obtain a suitable sample of the material for analysis in the device. This is often impractical when considering natural scenes, where the object under consideration cannot be moved out of the scene and into the laboratory. There has, therefore, been a demand for cheaper and more convenient methods of measuring surface reflectance. Of particular interest has been the use of images, which are themselves a representation of surface reflection, to determine the reflection characteristics of observed surfaces.

Research into this subject area has been led mainly by two activities: computer graphics and machine vision. The computer graphics community has required the accurate modelling of materials for environment simulation. The apparent realism of a virtual environment is enhanced through the accurate modelling of light's interaction with the scene. The use of precise geometric scene modelling may be wasted if the reflection models used to render the scene are simplistic, resulting in apparently unnatural images. Traditionally, virtual environment builders have empirically estimated surface characteristics but the increasing demand for more physically-based

¹See Chapter 3 and Appendix A for an explanation of the radiometric terms used in this chapter.

lighting simulation has led to the need for accurate BRDF measurements for a wide variety of materials. Many reflection models have been proposed by the graphics community (see Chapter 3), some derived from a physical consideration of the processes involved, but it is the parameters of these models, for a variety of different surface materials, that are required. An inexpensive and accessible method of obtaining parameter values is therefore highly desirable.

Machine vision has also been a driving force behind image-based measurement of surface reflection. Such data can be used for a variety of tasks. The segmentation of images into contiguous regions, representing the projection of a scene object, has traditionally relied on the detection of high contrast areas to locate object boundaries. However, such contrast may be due not only to object changes but also to reflection effects such as specular highlights. Object recognition may be aided by the use of reflection characteristics in addition to shape [66]. Observed surface reflection can also be used to aid product inspection for quality control [82]. However, most work in this field has been associated with photoclinometric or *shape-from-shading* techniques [50]. Such techniques attempt to recover scene geometry from images based upon the variation in observed shading. However, the interpretation of image shading requires knowledge of the surface reflection model. Therefore, much work on reflection model estimation has been undertaken in the context of shape-from-shading and its related topics.

This chapter presents work related to the estimation of a reflection model for a surface observed in image data. It initially presents methods for the direct measurement of surface BRDF which may only be obtained by sampling the full hemisphere of possible illumination and viewing directions. It will then discuss intensity-based approaches using a limited number of scene observations to estimate surface reflection characteristics. Much of this work is related to the shape-from-shading problem. It will then proceed to discuss colour reflection analysis which has, in general, been associated with image segmentation and object recognition.

2.1 Image Based Measurement of BRDF

A gonireflectometer measures surface reflectance characteristics by sampling the full hemisphere of possible incident and reflection directions. This is achieved through accurate manipulation of the sample material, the light source, and the sensor device. This section presents two image based methods which, like the gonireflectometer, consider the full range of geometries within

which reflection can take place. In this respect this section differs from subsequent sections, which attempt to determine surface reflectance from a restricted number of images. The accuracy of the resulting reflectance models is determined by the density with which measurements are taken.

Ward [141] recognised the need for an inexpensive device for measuring reflectance and designed an image-based gonireflectometer to measure the anisotropic reflectance of a material sample. The *imaging gonireflectometer* uses a half-silvered hemisphere and a fish-eye lens to capture the hemispherical reflectance from a surface in a single image. In this way the reflected radiance in the hemisphere of possible reflection directions, due to a single incident direction, can be obtained. The system measures the reflected radiance relative to a standard Lambertian sample of known reflectance, to aid calibration. The apparatus consists of a white collimated light source, outside of a half-silvered hemisphere, which illuminates the surface sample within. A static camera is used to record the image obtained from the reflective surface of the hemisphere. A number of images are obtained as the light source and the material sample are manipulated. Ward obtained reflectance measurements for a wide variety of materials including brass, aluminum, wood, cardboard, ceramic, plastic, paper and paint. However, his system is unable to measure ideal, or near ideal, reflectance from specular surfaces due to the limited directional accuracy. The resulting measurements have been used to derive a model of surface reflectance that is applicable to a wide variety of surface types.

Dana et al. [16] used a similar, image-based, approach to measure anisotropic surface reflectance from a variety of material samples. They used a radiometrically calibrated camera and extended light source to measure reflected radiance. The resulting measurements were fitted, using a least-squares approach, to both the Oren-Nayer [102] and Koenderink [68] reflection models. Both camera and light source were manipulated about the surface sample to record reflectance from the full range of incident and reflection directions. Up to 200 geometry combinations were considered in this way. From the resulting measurements a database of colour reflection model parameters has been produced. This includes data for 61 material samples such as frosted glass, feathers, cotton, cardboard, wood, orange peel and straw.

Whilst these approaches to image-based reflectance measurement are relatively inexpensive, they have not addressed the issue of practicality. They are, like the gonireflectometer, laboratory-based systems which require a suitable material sample to be obtained. This is often an impractical requirement. However, such studies do allow for the evaluation of reflection models and the accu-

mulation of material data. They also serve to show that radiometric measurements can be obtained using simple imaging devices, such as CCD cameras, provided suitable calibration is performed (see Chapter 5).

2.2 Greyscale Image Analysis

This section considers work on estimating surface reflection properties from a restricted number of greyscale images. Typically, these images view complex scenes and not simply individual objects. In these cases it is not possible to directly measure BRDF, since the full range of illumination and viewing geometry cannot be sampled in a single, or a limited number of images. Therefore, these approaches attempt to find the parameters which yield a *best fit* between observed data and a given model of reflection. Other methods attempt to detect, rather than measure, features such as interreflection or specular reflections, to aid image segmentation or other machine vision tasks.

As stated above, related work in this area of study has typically been associated with other image processing tasks such as shape-from-shading or image segmentation. Little work has been performed on the task presented in this thesis of estimating reflection model parameters from a complex scene of known geometry. This is probably due to the relatively recent advances in range finder technology, which have meant that inexpensive devices can now be obtained to accurately determine scene geometry. Such devices can be used to obtain 3D scene geometry (or relative location and surface normal) irrespective of surface material or illumination. The resulting geometric data may then be used to aid the analysis of observed image shading. However, shape-from-shading methods are still required for tasks such as remote sensing where the terrain is not known *a priori*.

2.2.1 Single Image Methods

The problem of *shape-from-shading* (SFS) from a single image was originally presented by Horn [49]. Horn's original work showed that a *reflectance map* can be used to relate image shading to surface orientation independently of position within the image. In its most general form, Horn's method is able to determine the shape of surfaces, with arbitrary reflectance, under perspective projection and illumination from a nearby light source. A first-order partial differential equation

known as the *image brightness equation* is derived in terms of five dependent variables. These define the surface patch location and gradient in a viewer-oriented co-ordinate space. A set of five ordinary partial differential equations are then derived and solved using the characteristic strip method which places constraints upon the resulting surface topology. However, this method suffers from several practical problems including noise sensitivity and error accumulation in numerical integration of the differential equations [75].

The original SFS problem may be simplified by considering orthographic projection, a distant light source from a known direction, and Lambertian surfaces of known constant albedo. These assumptions enable the reflectance map to be defined in terms of two variables, owing to the view independence of the Lambertian reflection model [51]. However, such precise knowledge of the environment severely limits the applicability of the method. There has, therefore, been much work on improving the original SFS algorithm to relax some of these assumptions. Light source position, for example, may be estimated by a number of methods [28, 104]. Of particular interest here are approaches which estimate the Lambertian surface albedo in addition to shape [153]. Horn and Brooks [50] and Zhang et al. [152] provide a thorough performance analysis of the major SFS algorithms.

Natural scenes typically deviate from the Lambertian ideal and contain rough surfaces which exhibit both diffuse and specular reflections. The reflection characteristics of such surfaces cause difficulties for traditional SFS algorithms. There is therefore a need for methods that are applicable to a wide variety of surface types. This requires the definition of reflectance maps to model the reflection from differing materials. Various improvements over the Lambertian reflection model have been proposed [102, 127, 150]. These attempt to account for attributes such as surface roughness and glossy reflections. Healey and Binford [43] show that a Torrance-Sparrow reflection model [134] may be used to determine shape from rough surface specular reflections. A combination of the Lambertian and Torrance-Sparrow reflection models, in the context of the SFS problem, has also been considered [136]. A recent approach to SFS using a generalised reflectance map is given by Lee and Kuo [75]. Typically, the resulting estimation of surface topology leads to the evaluation of parameters for the applied reflection model.

Interreflection, or mutual reflection, between surfaces in a scene also presents problems for traditional SFS algorithms. Koenderink and van Doorne [67] have presented a formal treatment of the process of interreflection between Lambertian surfaces of arbitrary shape and varying albedo.

Nayer et al. [92] have looked at the analysis of such interreflections within an image as indicators of surface shape and albedo. They again assume Lambertian surfaces, but present an iterative approach that recovers both surface shape and albedo.

The SFS problem has also been considered in the case of more natural illumination. Langer and Zucker [73] have approached the problem of SFS under a diffuse hemispherical illuminant which, they suggest, is similar to that of an overcast sky. They show that shape recovery is possible but assume the scene to be composed of Lambertian surfaces of known albedo.

The *a priori* knowledge of scene geometry aids the estimation of surface reflectance properties. Karner et al. [58] used an image-based system to estimate the parameters of the Ward reflection model [141] for a planar material sample observed in a precise position. They used a diffuse reflectance standard within the same image for calibration and comparison purposes. Two measurements are obtained, with and without a fixed point light source, to compensate for ambient and stray light. A ratio of the known sample and the measured sample image intensities across the image are used to derive the surface BRDF parameters. Results are shown for both metallic and dielectric materials.

Ikeuchi and Sato [53] have studied the fusion of a single range and intensity image to obtain object shape and surface reflectance. The proposed reflection model has three components: a diffuse Lambertian lobe; a specular lobe conforming to the Torrance-Sparrow rough surface model [134]; and a specular spike modelled using a delta function. Using least-squares fitting, they estimated model parameters for an observed object. Results are shown for Lambertian and specular objects assuming uniform reflectance properties over the whole surface.

In the field of remote sensing the shape-from-shading problem is more commonly termed *photoclinometry*. Though the images considered here are typically obtained at great distance, such as those from a satellite, aerial observation, or a planetary image, the problem is very similar to that of SFS. In this case, parameter values are estimated assuming a given surface geometry. The resulting values provide information concerning the state of the surface, such as particle size or undulation. The validity of the results is determined by how well the reflection model characterises the observed surface. Hapke's extensive reflection model [37, 38] has been fitted to observed shading sensed from a number of surface types including soil [109], snow and ice [11, 140]. Other models have also been fitted to observed data for surfaces such as forest canopy [33], sand and

cloud [115].

2.2.2 Multiple Image Methods

Reflection model estimation, and shape recovery, using a single intensity image is highly under-constrained due to the limited variation in illumination geometry that is represented. The degree of geometry variation is dependent upon the change of surface normal across the observed object. The use of further images can help to constrain the problem of reflection model estimation. Two approaches may be used: observing a static scene under varying illumination conditions, or observing a moving object. This section considers such approaches. As with the single image methods presented above, much of the work presented here is primarily concerned with the estimation of surface shape.

Static Scenes

Woodham introduced the concept of *photometric stereo* whereby a scene is illuminated from different, known light source positions [151]. Using Horn's reflectance map, the observed image intensities, in each of the images, can be used to determine surface shape. Woodham showed that, in the case of a Lambertian surface, three images, illuminated by non-coplanar light sources, are required to sufficiently constrain the problem of shape estimation. The resulting inversion of the image formation process allows surface albedo to be estimated. The same approach can be used for the analysis of ideal specular surfaces illuminated by extended sources [52]. Coleman and Jain [13] have extended the principle to analyse textured and specular surfaces. In this case it is shown that a four light photometric stereo technique is necessary to extract the shape since specular highlights may occur in any one of the three images used in Woodham's method. This is true provided that specular highlights do not overlap between images. Although no attempt is made to determine the level of specular reflection, they are able to subtract it from the images to provide intrinsic Lambertian images for analysis. A four-light illumination has also been used by Solomon and Ikeuchi [123] to determine the roughness of a specularly reflecting Torrance-Sparrow [134] type surface. Silver [122] has also developed ways of applying the basic photometric stereo method to surfaces of differing reflectance properties. This work also showed that photometric stereo could be conducted using experimentally measured reflectance maps.

The similar concept of *photometric sampling* has been introduced by Nayer et al. [93]. In this case extended, rather than point sources, are used to illuminate the object from a variety of source positions. The resulting sequence of images is used for shape and reflectance estimation. A least-squares technique is used to estimate the parameters of a combined Lambertian and simplified Torrance-Sparrow model [134], though experiments focus on the analysis of smooth surfaces. Again, it is assumed that highlights do not overlap between images. Though quantitative results are not presented, the subtraction of the specular component to aid shape recovery is applied and the relative strength of the diffuse component is calculated. Kay and Caelli [60] applied the same reflection model and used non-linear regression techniques for the analysis of glossy reflections from rough surfaces. It was shown that parameter estimation may be carried out locally at each pixel rather than globally on the entire image. In this case the light source positions are contained within a plane and rotated about the camera's principal axis. Tagare and deFigueiredo [126] have also considered the estimation of shape and reflectance for Lambertian and glossy specular objects, though they are unable to provide a stable solution for reflection model parameter estimation.

Grimson has shown that existing photometric stereo methods can be modified to be used for binocular stereo where the scene is observed from two differing camera positions [34]. Grimson showed that surface orientation can be estimated along with reflectance. Again, Horn's reflectance map is used; however, in this case reflectance is modelled using Phong's empirical model [108].

The use of data obtained by a range finder enables surface normal estimation to be performed independently of model parameter fitting. Kay and Caelli [59] fuse range and brightness images to derive surface shape and reflectance. In this case the reflection model parameters are estimated locally for each image pixel. Brightness images are obtained under 4 or 8 differing illumination conditions. However, this method is unable to cope with highly textured surfaces or scenes consisting of many varied materials.

The separation of specular reflection highlights from images can be achieved through the use of polarising filters. Wolff and Boult present a unified framework for the analysis of polarised images and present the *Fresnel reflectance model* [149]. Using this model to interpret differently polarised images enables the classification of smooth metallic and dielectric materials [148]. This is due to the differing ratio of Fresnel coefficients for conductive and dielectric materials. Wolff also considered reflection from rough surfaces in accordance with the Torrance-Sparrow model [134] and successfully determined quantitative diffuse and specular reflection components [147].

Moving Scenes

Scenes in which objects move provide an alternative to the photometric stereo and sampling methods presented above. In this case the motion, and hence changing illumination and viewing geometry, is used to constrain traditional SFS methods. Pentland uses optic flow and *photometric motion* for the estimation of surface shape and Lambertian albedo [105]. Again, Horn's reflectance map is used to relate image intensity to surface orientation, but in this case a sequence of images observing a moving object under a known illuminant is considered. This requirement for *a priori* knowledge of the illuminant has been recently examined by Mukawa [88].

Lu and Little [79] present a method which observes a rotating object illuminated by a collinear light source. The light source lies on or near the optical axis of the camera. It is shown that surface reflectance can be directly estimated from the image sequence which observes a complete rotation of the object. The reflectance model is estimated using *singular surface points* whose normals are in the viewing direction and as such are assumed to exhibit maximum image intensity. Given the 3D location of these singular points, and their brightness values in the image sequence, the Lambertian reflectance of the surface can be estimated.

2.3 Colour Image Analysis

Surface colour is an important characteristic of an object. However, apparent colour is a product of both the spectral nature of the illumination and material properties. A given object may appear to have very different colouring when viewed in illumination of differing spectral composition. *Colour constancy* is the ability to recover a surface description of colour which is independent of the illumination. Therefore, even if illumination colour changes, the surface descriptor of colour remains constant. This provides a useful attribute for object recognition. Approximate colour constancy is exhibited by the human visual system. Humans perceive the colour of a surface as an invariant characteristic. For example, a human observer recognises grass as being green despite the changing spectral composition of daylight illumination. Conversely, a camera would give a high red channel response, rather than green, to grass observed at sunset.

Section 2.3.1 presents methods of colour constancy that are able to determine the spectral properties of a material independently of the illuminant. Typically, such methods assume a scene

of known geometry and Lambertian materials. Colour constancy algorithms typically neglect rough surfaces and glossy reflections. Attempts to analyse more complex reflection characteristics have generally focused on the use of the dichromatic reflection model [119]. These methods are presented in Section 2.3.2. Finally, the use of alternative models is considered.

2.3.1 Colour Constancy

The first computational model of human colour constancy was proposed by Land and McCaen [72]. Their *retinex theory* assumes a Mondrian² world which consists of planar patches of differently coloured paper. The illumination across this Mondrian world is assumed to be smoothly varying over the observed scene. As such, sharp changes in colour signal intensity can be attributed to object boundaries, whereas smooth changes are due to illumination variation. Brainard and Wandell [9] have examined the properties of the retinex algorithm and have found it to be an effective colour constancy method. In general the algorithm can determine constant colour descriptors despite changes in illumination. However, if the scene surrounding a patch is changed, different colour descriptors are generated.

More recent approaches to colour constancy have used a finite-dimension linear model in which surface reflectance and illumination are both expressed as a weighted sum of fixed basis functions [10, 21, 25, 47, 80]. The task of colour constancy, therefore, becomes that of estimating the reflectivity weights for the object and the illumination weights. Typically the scene is assumed to be Mondrian and composed of Lambertian surfaces. Maloney and Wandell [80] defined a number of conditions to be satisfied for efficient colour constancy. They show that illumination must be constant over a given segment of an image and a sufficiency of different chromatic information must be available. They also show that, if n weighted basis functions are used as a surface descriptor, then at least $n + 1$ spectral samples of reflected radiance are required. This requirement has since been improved upon by Ho et al. [47] who show three parameter recovery of a surface descriptor and illumination descriptor from a three channel colour signal. Recently, Finlayson et al. [23] have shown colour constancy using two reflectance measurements under spectrally distinct illuminations. Interreflection between surfaces has also been considered in the context of colour constancy [25].

²After the Dutch abstract artist Piet Cornelius Mondrian.

Supervised colour constancy uses objects of known properties as reference materials within the same image. From the reflected signal from these known samples the illuminant colour can be estimated [98]. The resulting illuminant information can be used as a constraint in determining object colour. A similar approach is proposed by Ohta and Hayashi for daylight illuminated scenes [100]. Knowledge of the possible spectral composition of daylight illumination allows the performance of existing colour constancy algorithms to be improved.

The extension of colour constancy to more natural scenes, which have varying scene geometry and surfaces which exhibit glossy reflection, has been considered by D’Zmura and Lennie [21]. They assumed that the specular component of reflectance is constant with wavelength and show that the hue of a surface is constant with respect to changing geometry. Tominaga and Wandell [133] also considered scenes which have a spectral reflection component and varying geometry. They employed the dichromatic reflection model [119] to describe interface and body reflection processes, typically observed from plastics.

Tominaga has recently presented a method for the measurement of surface spectral reflectance from a scene composed of objects which exhibit highlights [131]. In this case a monochrome CCD camera, with extended dynamic range, is used with filters so that a six channel image is recorded of the scene. The camera’s sensitivity to each filtered wavelength is determined by prior calibration. Results from experiments on a small number of surface samples indicate that accurate estimation of surface colour can be achieved in addition to the illumination’s spectral composition.

2.3.2 Colour Model Estimation

Whilst colour constancy attempts to determine spectral reflectance properties, a full colour description of a surface material must also consider surface roughness and glossy reflections. Such reflection models are necessary for the analysis and interpretation of natural scenes which include a wide variety of material types. Much of the work in this area of study has focused on the use of the dichromatic reflection model [119], however alternative models have been used for the analysis of colour reflection.

The Dichromatic Reflection Model

The dichromatic reflection model, originally proposed by Shafer [119], provides a model of reflection for inhomogeneous dielectric materials. In particular, it considers smooth, pigmented materials which contain colouring dye particles suspended in a medium. This class of materials includes plastics, paint, ceramics and some textiles. Only two lobes of reflection are considered: ideal, specular interface reflection from the material surface; and diffuse body reflection due to scattering within the bulk of the material. The model is typically used to aid segmentation of colour images through the ability to detect and remove specular reflections which are assumed not to have been coloured by the pigment particles [66]. Shafer has shown that RGB pixel values corresponding to a single material will lie on a parallelogram within the colour space [119]. In a similar approach, H.-C. Lee et al. [74] have developed the Neutral Interface Reflection (NIR) model which has been shown to be suitable for identifying specular reflections.

Klinker [66] and Gershon [27] independently showed that dichromatic reflection from inhomogeneous dielectric materials is characterised by two connected clusters in RGB colour space. Klinker describes the resulting shape as a *skewed-T* formation. From the orientation of these clusters, the illuminant colour and material pigment can be determined, as can the level of specular reflection. This information may be used for segmentation purposes or for the estimation of material colour and the relative contributions of body and interface reflection. Hashimoto et al. [39] have used this characteristic clustering of pixel values in colour space for the recognition of material types under white light. They applied the dichromatic model to a single image to distinguish between metallic, matte non-metallic, and glossy non-metallic surface types. However, they made no attempt to quantify the specular surface reflection characteristics. It has also been noted that the characteristic clustering occurs in other colour spaces, such as the HSI space [1, 76]. The conversion of image data to these spaces provides a more direct correspondence between clusters and surface shading due to body and interface reflections.

Tominaga and Wandell [132, 129] have tested the adequacy of the dichromatic model for characterising reflection from inhomogeneous materials. It is reported to apply well for plastic, paint, ceramic, vinyl, fruit and leaves; however, it failed on metal, textiles and paper. Tominaga [130] has extended the standard model with the inclusion of a wavelength dependent interface reflection component. In the case of metallic surfaces, the diffuse body component is

neglected. They show that three variants of the basic dichromatic model are suitable to describe the reflectances of most materials.

As with greyscale images, interreflection between surfaces can cause problems for image analysis methods. Bajcsy et al. [1] and Novak and Shafer [97] have used the dichromatic model to detect highlights and interreflections in images. This has been performed as an aid to image segmentation where the characteristic colour bleeding presents difficulties. It is noted that extra clusters in colour space are observed due to interreflection processes.

Real scenes are complex, with objects that may be textured or have patches of different reflectance properties. Cluster analysis, in the context of the dichromatic model, assumes smooth surfaces and an object which has been sufficiently segmented from other image data. These conditions allow surface colour and illuminant analysis to be estimated from the cluster orientations. Rough surfaces and image noise cause the clusters to spread and hence make it difficult to separate the reflection into its components. This problem is highlighted by Novak and Shafer [97] and S. W. Lee [76]. Novak also showed that histogram shape is related to illumination intensity and the phase angle between camera and light source.

Sato and Ikeuchi [116] have used the cluster analysis method in the context of an image sequence. In this case the object is observed under an extended light source which is placed in a number of differing positions, as with the photometric sampling method [93]. The resulting pixel values may be plotted in a five dimensional *temporal-colour space*. It is shown that clustering also occurs in this space. From analysis of the clusters, the orientation and the reflectance can be determined for each individual pixel. This is a departure from the previous cluster based methods which have required pixel values from a region of the object. The method cannot be applied where interreflection occurs, though the method is shown to be applicable to both dielectric and metallic surfaces.

Sato and Ikeuchi [117] have extended the work presented in [116] to scenes observed under daylight illumination conditions. Again, the dichromatic model is used to interpret a sequence of images obtained over a period of time. The duration of the image sequence is such that the movement of the sun provides an illuminant suitable for analysis by the photometric sampling method. In this case, the natural daylight illumination is modelled by a sun, represented as a Gaussian disc, and a sky which is assumed to have uniform radiance over the illuminating hemisphere. They

show that surface normal recovery is under-constrained due to the sun's limited trajectory across the sky hemisphere. Using the dichromatic model, they analyse a real image sequence and estimate the colour vectors due to Lambertian and specular reflection. Analysis focuses upon the detection of highlights due to the sun and the subsequent separation of reflection from sun and sky illumination. This is shown to allow shape determination rather than quantitative reflection model parameter estimation.

In a similar approach to Wolff [149], Nayer et al. [91] have recently used polarisation for the analysis of colour images. In the context of the dichromatic reflection model, polarisation can be used to separate specular and interface components. In this case, it is assumed that light becomes polarised after having been reflected from the specular interface of the material. This allows for specular component removal from the images to aid analysis using the resulting diffuse component images. The algorithm is applied to complex scenes which include textured objects and interreflections. Results show that estimates can be made of the specular component.

Alternative Colour Models

Owing to the many successful applications of the dichromatic reflection model, and its derivatives, little work has been performed on colour image reflection analysis using alternative reflection models. This section presents two alternative approaches which have utilised range finder data to aid scene analysis.

Baribeau et al. [2] have used range and colour images to estimate local reflection properties in the context of the Torrance-Sparrow model [134]. The use of this model allows the analysis of rough surfaces which exhibit glossy reflections. An extended source is used to illuminate the scene to enable surface roughness and Fresnel reflectance parameters from three selected object areas to be estimated.

Sato and Ikeuchi [118] have also employed a laser range finder to estimate object shape. In this case, however, the object is observed whilst it performs a 360° rotation, with the illuminant position similar to that of the camera. The resulting sequence of range images is used to generate a set of triangular patches to represent object shape. The colour image sequence, of up to 120 images, is then separated into diffuse and specular reflection components using the method presented in [116]. However, in this case the reflection model is a linear combination of the Lambertian and

a simplified Torrance-Sparrow model [134]. This allows surface roughness to be accounted for in the specular component. Results are shown for a glossy plastic object with parameters estimated using Levenberg-Marquardt regression [5].

2.4 Summary

This chapter has presented a review of work related to reflection model estimation from image data. It has shown that imaging devices, such as CCD cameras, may be successfully used for the measurement of reflection provided that suitable calibration is undertaken. The use of imaging devices allows more complex and natural objects to be analysed within their environment. This method of analysis is to be contrasted with the use of specialised devices, such as the gonio-reflectorimeter, which require a material sample to be obtained for laboratory-based experiments.

The estimation of surface reflection properties from image data has typically been performed in conjunction with a specific task, such as image segmentation, object recognition or shape estimation. These methods often only attempt to detect reflection characteristics, such as highlights or interreflection, that may assist in the primary objective. These methods result in a qualitative analysis of surface reflection rather than an estimation of reflection model parameters. Such results may be used, however, to provide a broad classification of observed objects by their material type.

The measurement of spectral material properties presents many difficulties for image based systems. This is due to the observed reflection being a product of both the illuminant and material properties. There has, therefore, been a need for an illuminant invariant description that can be recovered from image data. This problem has been addressed by colour constancy methods. However, as this chapter has shown, these typically require unnatural constraints to be placed on the observed scene.

The study of natural scenes, illuminated by natural illumination, has received little attention. Those studies that have examined machine vision tasks under such complex illumination have typically employed simplistic reflection models which are applicable only to specific material types. Quantitative analysis of reflection in such environments, accounting for surface roughness and a variety of materials, has not been attempted.

Chapter 3

Surface Reflection

Our perception of objects is governed by the interaction of light with the environment. The physics of light propagation and reflection has been the subject of intense study for hundreds of years. Today the true nature of light still remains unresolved. Two theories have come to the fore, each with its protagonists. James Clark Maxwell (1864) formalised the view of classical physics by defining the relationship between electric and magnetic fields. Maxwell determined that electromagnetic waves propagate with the speed of light indicating that light itself is an electromagnetic wave [121]. More recently the quantum model of light has found favour, depicting light as a distribution of charge packets. Neither the classical wave model nor the quantum model is able to adequately explain the phenomena that light exhibits in nature [124]. This chapter will focus on the wave model which is able to describe most of the reflection characteristics exhibited by natural surfaces.

The study of light's interaction with the environment needs to encompass the three processes that light may undergo on meeting a material interface: *reflection*, *absorption* and *transmission*. This chapter focuses on the study of opaque objects which do not transmit light. As shall be shown in the proceeding sections, a complete study of reflection must not neglect light transmission. As Maxwell's equations show, light reflected from opaque conductive materials may be attenuated by the transmissive nature of the material. In the case of dielectric surfaces this chapter will concentrate on the processes of reflection and absorption only. This chapter will also assume an air-object interface which is that most commonly encountered in natural scenes.

Maxwell's equations are used to describe the interaction of electromagnetic waves with any

isotropic medium under the condition of no accumulation of static charge. For the processes of visible light interaction, this chapter considers only radiation in the visible region of the spectrum (380-770 nm). Maxwell's equations cover wave-surface interaction for ideal surfaces which are optically smooth, clean and reflect light in an ideal specular fashion. In this context a surface is deemed smooth if surface irregularities are small compared to the wavelength of incident light (see Section 3.3). These equations provide a basic understanding of light reflection but are complex to compute. They also depart from observed measurements due to surface roughness, surface contamination, material impurities and crystal structure modification by surface working. It is these natural surface characteristics that must be addressed by a practical model of surface reflection.

This chapter presents models of light reflection that are applicable to this thesis. These models have been developed from the study of radiant heat transfer and applied to computer graphics and machine vision tasks. Any model of surface reflection presents either a solution or an approximation to Maxwell's fundamental equations. This chapter begins by defining a method of measuring and describing a surface's reflecting characteristics. Fresnel's solution to Maxwell's equations for ideal surfaces will then be considered. This solution provides the basis for the computational models which follow. Models of surface imperfections are also presented as these enable the theoretically ideal reflection models to be applied to real-world surfaces.

3.1 Reflectance: The Measure of Reflection




Figure 3.1: *Surface reflection geometry.*

Reflection is the process by which incident flux leaves a stationary¹ surface from the incident side without change in frequency [95]. Figure 3.1 shows the reflection geometry for a differential surface element dA . Note that the angles of incidence and reflection, θ_i and θ_r , are measured relative to the surface normal, n . The azimuthal angles, ϕ_i and ϕ_r , are measured relative to a reference vector perpendicular to the surface normal.

A surface may not reflect the same quantity of flux for each possible incident direction. It is also possible that the surface may cause scattering of the incident light into a distribution of reflection directions. This section presents distribution functions which enable the reflection properties of a surface to be fully described by a single function.

3.1.1 Reflectance

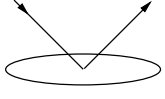
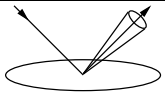
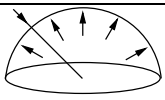
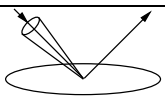
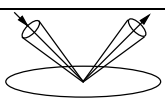
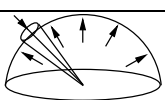
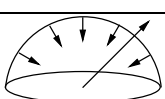
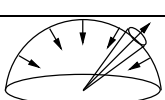
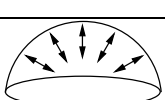
Reflectance, ρ , represents the ratio of reflected flux to incident flux² [31]. It is a function of the geometry of the incident and reflected flux and may be dependent upon the wavelength of incident light. When considering flux incident on or reflected from a surface it is necessary to define the distribution of directions within which the flux is constrained. Different reflectance terms are used to specify the geometry under consideration. Table 3.1 shows the nine possible reflectance geometries for a planar surface element. In this context *directional* refers to flux within a differential solid angle ($\vec{\omega}$), *conical* refers to flux within a cone of finite solid angle ($\Delta\vec{\omega}$) and *hemispherical* refers to flux within a hemisphere of directions ($\omega = 2\pi$). The directional-hemispherical reflectance is also referred to as the surface *albedo*.

Reflectance defines the total quantity of flux reflected by a surface. The hemispherical and conical reflectances give no indication as to the relative distribution of the flux, either incident on or reflected from a surface. To consider the spatial distribution of flux a more fundamental measure of reflection is required.

¹If the surface were to be in motion the reflected light would be subject to Doppler shift.

²See Appendix A for an explanation of the radiometric quantities used in this chapter.

Table 3.1: The nine reflectance geometries.

Bi-directional		$\rho(\vec{\omega}_i \rightarrow \vec{\omega}_r)$
Directional-conical		$\rho(\vec{\omega}_i \rightarrow \Delta\vec{\omega}_r)$
Directional-hemispherical		$\rho(\vec{\omega}_i \rightarrow 2\pi)$
Conical-directional		$\rho(\Delta\vec{\omega}_i \rightarrow \vec{\omega}_r)$
Bi-conical		$\rho(\Delta\vec{\omega}_i \rightarrow \Delta\vec{\omega}_r)$
Conical-hemispherical		$\rho(\Delta\vec{\omega}_i \rightarrow 2\pi)$
Hemispherical-directional		$\rho(2\pi \rightarrow \vec{\omega}_r)$
Hemispherical-conical		$\rho(2\pi \rightarrow \Delta\vec{\omega}_r)$
Bi-hemispherical		$\rho(2\pi \rightarrow 2\pi)$

3.1.2 The Bidirectional Reflectance Distribution Function

Nicodemus et al. [95] presented a defining paper which provided a standard framework for the description of light reflection from a surface. The fundamental contribution of this work was the presentation of the *bidirectional reflectance-distribution function* (BRDF). This provides a flexible and general mathematical function with which to describe the anisotropic reflection of incident flux from most surface types [31]. The BRDF, f_r , describes the distribution of reflected light as a function of the incoming and outgoing directions and relates reflected radiance, L_r , to differential incident irradiance, L_i ;

$$f_r(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{L_r(\theta_r, \phi_r)}{L_i(\theta_i, \phi_i) \cos \theta_i d\omega_i} \quad [\text{sr}^{-1}]. \quad (3.1)$$

The BRDF has the following properties:

1. The BRDF obeys the *Helmholtz reciprocity principle* so that if the incident and reflection directions are reversed then f_r is unchanged,

$$f_r(\vec{\omega}_i \rightarrow \vec{\omega}_r) = f_r(\vec{\omega}_r \rightarrow \vec{\omega}_i). \quad (3.2)$$

2. The BRDF is, in general, *anisotropic*. That is, if the surface is rotated about the surface normal whilst the incident and reflected directions remain unchanged then the value of f_r may change. A surface which is *isotropic* has a BRDF which is independent of the surface orientation,

$$f_r(\theta_i, \phi_i + \phi; \theta_r, \phi_r + \phi) = f_r(\theta_i, \phi_i; \theta_r, \phi_r). \quad (3.3)$$

Nicodemus et al. show that each of the nine reflectance functions shown in Table 3.1 can be defined using the BRDF [95]. The BRDF can be readily extended to include a dependence on the wavelength of the light under consideration. With this extension the BRDF can be used to describe wavelength attenuation by the surface material. This is the cause of apparent object colour. A wavelength dependent BRDF can be defined as,

$$f_r(\theta_i, \phi_i; \theta_r, \phi_r; \lambda) = \frac{L_r(\theta_r, \phi_r; \lambda)}{L_i(\theta_i, \phi_i; \lambda) \cos \theta_i d\omega_i}. \quad (3.4)$$

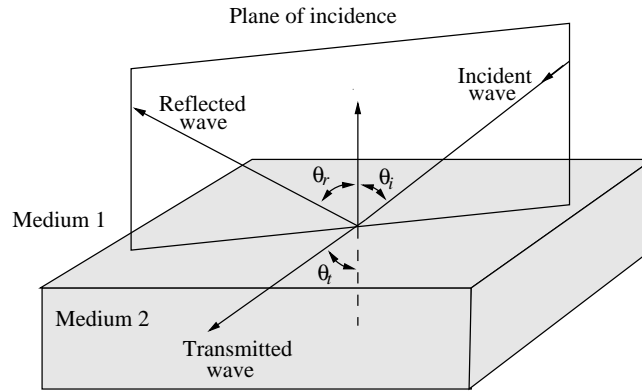


Figure 3.2: Behaviour of an incident wave at the interface between two media.

3.2 Fresnel's Laws of Reflection

Maxwell's equations provide the basic *law of reflection* which state that incident and reflected waves have directions symmetrical with respect to the normal at the point of incidence and are contained within the plane of incidence, $\theta_i = \theta_r$, (see Figure 3.2). Such ideal reflection is termed *mirror* or *specular* reflection and the surface is termed a *specular surface*.

Maxwell's equations also provides the derivation of Snell's law of refraction, again with the transmitted wave refracted within the plane of incidence,

$$\frac{\sin \theta_t}{\sin \theta_i} = \frac{n_1}{n_2}, \quad (3.5)$$

where n_1 is the refractive index of medium 1 in which the incident wave propagates and n_2 is the refractive index of medium 2. In the special case where medium 1 is air, $n_1 \approx 1$ and hence

$$n_2 = \frac{\sin \theta_i}{\sin \theta_t}. \quad (3.6)$$

The refractive index of a material is determined by the speed with which light propagates in the material relative to a vacuum, $n = c/c_m$, where c is the speed of light in a vacuum and c_m is the speed of light in the medium [31]. Since the refractive index is a function of wavelength it can be seen that the angle of refraction is dependent upon the wavelength of incident light. In the case of materials with finite conductivity (such as metals) the refractive indices become complex quantities,

$$\bar{n} = n + i\gamma, \quad (3.7)$$

where \bar{n} is the complex index of refraction and γ is the materials *absorption coefficient*.

Fresnel's coefficients [144] give the amplitude of the reflected electric field in terms of the electromagnetic fields parallel and perpendicular to the plane of incidence. From these the isotropic reflectance for unpolarized radiation can be determined,

$$\rho(\vec{\omega}_i \rightarrow 2\pi; \lambda) = \frac{1}{2} \frac{\sin^2(\theta_i - \theta_t)}{\sin^2(\theta_i + \theta_t)} \left(1 + \frac{\cos^2(\theta_i + \theta_t)}{\cos^2(\theta_i - \theta_t)} \right). \quad (3.8)$$

Equation 3.8 is known as Fresnel's equation. It defines the directional-hemispherical reflectance as a function of wavelength and the angle of incidence and refraction. Note that Fresnel's equation obeys Helmholtz's law of reciprocity. The equation is a minimum when $\theta_i = 0$ and is unity (no light is absorbed by the material) when $\theta_i = \frac{\pi}{2}$. Fresnel's equation shows that metals have a higher reflectance than dielectrics. For example silver and aluminum reflect over 90 per cent of all visible light [124], and the reflectance of metal varies considerably with wavelength.

Fresnel's laws of reflection characterise the reflected and refracted light as they depend on wavelength, polarisation and angle of incidence. Since they are derived directly from Maxwell's equations they assume a surface is optically smooth. Models of surface reflection and refraction build upon Fresnel's laws and extend them to encompass rough surfaces of differing material types.

3.3 Modelling Surface Structure

The preceding sections of this chapter have so far assumed an ideal reflecting surface. That is, a surface which is smooth, clean and reflects light in a specular fashion. Such surfaces are rarely encountered in the natural environment and depart from this ideal in a number of ways.

The Rayleigh criterion is an approximate quantification that defines when a surface may be considered rough, or when specular reflections become scattered [144]. It states that a surface is considered rough when the change in surface height is such that reflected light undergoes interference. This interference occurs when light waves become out of phase due to the different reflection path lengths caused by the changes in height of a rough surface (see Figure 3.3).

The scattering of light from rough surfaces has been the subject of much study, especially in the field of photometry. These models have in general focused upon the particulate surfaces that are to be found on planetary bodies. The models thus derived [11, 37] may be applied to terrestrial

surfaces such as sand and dry soil. Stochastic models have also proved to be suitable for describing surface roughness, either through the use of fractals [20, 114] or Monte Carlo processes [145].

This section only considers isotropic surfaces and two popular models are presented. Such surfaces exhibit the same surface texture in all directions. An anisotropic surface may appear to have different surface textures when rotated about the surface normal [110, 141]. Whilst anisotropic surfaces are commonly encountered in natural scenes they have not been included in the initial investigations. Dirt and other surface impurities are also present in natural scenes. These effects have not been the subject of much study and are hence not considered here.




Figure 3.3: *Difference in path lengths for light reflected from a rough surface.*

3.3.1 Surface Height Distribution




Figure 3.4: *Surface height distribution model.*

The height of a point on a surface can be described by a stationary random function which has zero mean and is spatially isotropic. A Gaussian process is the common method for modelling surface height z , which is given by the probability distribution:

$$p(z) = \frac{1}{\sigma_z \sqrt{2\pi}} e^{-\frac{z^2}{2\sigma_z^2}} \quad (3.9)$$

where σ_z is the *rms roughness* of the surface. This, however, does not provide a full description of the surface geometry as there is no indication of the distance between surface peaks. Equation 3.9 can be extended to become a two point probability function which represents the two point height

distribution of the surface [42]:

$$p(z_1, z_2) = \frac{\exp\left(-\left(z_1^2 + z_2^2 - 2C(r)z_1z_2\right)/2\sigma_z^2(1 - C(r)^2)\right)}{2\pi\sigma_z^2\sqrt{1 - C(r)^2}} \quad (3.10)$$

where r is the horizontal Euclidean distance between the points and $C(r)$ is the *autocorrelation coefficient*. Figure 3.4 shows the profile of a surface which can be modelled using such a function. The autocorrelation coefficient is a circularly symmetric function of r ,

$$C(r) = e^{-\frac{r^2}{\tau^2}}, \quad (3.11)$$

where τ is the *correlation distance* which represents the average peak to valley distance [36]. Using such a representation, surface roughness can be modelled using just two parameters, σ_z and τ . The average slope of the surface facets for this representation is $\frac{2\sigma_z}{\tau}$ [4].

3.3.2 Facet Slope Distribution




Figure 3.5: *Facet slope distribution model.*

A popular alternative to the height distribution model considers the surface as a collection of planar microfacets. The surface has a mean normal and each microfacet has a local normal which deviates from the mean by an angle α (see Figure 3.5). The distribution of facet normals can be modelled by assuming α to be a random variable with a distribution rotationally symmetric with the mean surface normal. Assuming the mean surface normal to have a slope deviation of zero, the distribution of α can be produced using a Gaussian process [134],

$$p(\alpha) = ce^{-\frac{\alpha^2}{2\sigma_\alpha^2}}, \quad (3.12)$$

where σ_α is the *rms slope* and c is a normalisation constant. Cook [14] proposes using the Beckmann distribution, which has a similar shape but without the arbitrary constant³,

$$p(\alpha) = \frac{1}{4\sigma_\alpha^2 \cos^4 \alpha} e^{-(\tan \alpha / \sigma_\alpha)^2}. \quad (3.13)$$

³Cook's original paper [14] omits the 4 in the denominator. This was corrected by Hall [36].

This surface representation has advantages in its simplicity and can be used to represent most isotropic surfaces. It can, however, be difficult to visualise the true shape of the surface based on the single parameter σ_α . There is no direct comparison between the slope and height distribution models but He et al. suggest that rms slope is proportional to $\frac{\sigma_z}{\tau}$ [42]. Despite not providing a complete model of surface geometry the facet slope model is popular and the scattering of light from real surfaces has been found to be dependent on local slope rather than the height change of the surface [94].

3.3.3 Shadowing and Masking




Figure 3.6: *Facet shadowing and masking.*

Blinn [7] and Cook [14] introduced the effects of self-shadowing and self-masking by rough surfaces to the field of computer graphics. At large angles of incidence or reflection some surface facets may be shadowed and/or masked by other facets (see Figure 3.6). The net effect is to reduce the amount of reflection that would normally be observed from an ideal smooth surface. Blinn and Cook both adopt a geometrical approach to modelling the effects by assuming the surface facets to be arranged in symmetric V-grooves. Shadowing and masking has become an integral part of many other models of reflection [37, 42, 41, 101].

3.4 Modelling Surface Reflectance

The BRDF of a surface can be measured for a set of surface orientations and illumination geometries using a *gonioreflectometer* [89]. This device enables the surface reflectance characteristics to be measured and recorded in the form of a look up table. Use of the resulting data is, however, impractical as it may contain noise and will not cover the entire domain of the BRDF. It is therefore desirable to be able to represent the BRDF in terms of a functional model [31]. Such a model

would need to encapsulate all the features of the surface but should not be expensive in terms of computing time. Ideally, the model would have parameters with a physical meaning with which to describe the surface's reflection characteristics [36].




Figure 3.7: *Linear approximation to BRDF.*

Tagare and deFigueiredo [127] provide a good framework for the study of BRDFs and the different reflection mechanisms that contribute to them. They present the BRDF as being a linear composition of four approximating basis functions (see Figure 3.7). These represent an ideal specularly reflected ray, a foreshatter lobe, a normal lobe, and a backscatter lobe,

$$f_r(\vec{\omega}_i \rightarrow \vec{\omega}_r) = \kappa_{spec} f_{spec}(\vec{\omega}_i \rightarrow \vec{\omega}_r) + \kappa_{fsc} f_{fsc}(\vec{\omega}_i \rightarrow \vec{\omega}_r) + \kappa_{norm} f_{norm}(\vec{\omega}_i \rightarrow \vec{\omega}_r) + \kappa_{bsc} f_{bsc}(\vec{\omega}_i \rightarrow \vec{\omega}_r), \quad (3.14)$$

where κ_{spec} , κ_{frc} , κ_{norm} , κ_{bsc} are dimensionless constants which represent the contribution of each of the respective lobes. Each of these lobes may have a dependence on the wavelength of incident light and, hence, may attenuate the reflected wavelength. Few models employ all four terms, with the specular ray and backscatter terms often being omitted. This section presents the popular models used in computer vision and graphics with reference to these lobes of reflection.




Figure 3.8: *Light reflection from pigmented dielectric material.*

The dichromatic⁴ reflection model, introduced by Shafer [119], is an example of a lobed

⁴Similar to H.-C. Lee's *Neutral Interface Reflection* (NIR) model [74].

model. In this variant the backscatter and specular reflection components are ignored. Further to this, the glossy reflection component is assumed not to have been coloured by the surface material. The foreshatter BRDF approximation therefore has no dependence on the wavelength of incident light. The dichromatic model assumes the wavelength of normal lobe reflection to be determined by a combination of incident wavelength and selective material absorption. The glossy, interface, reflection is assumed not to have been attenuated and hence has the same wavelength as the incident light. Such a model is useful in describing the reflection from inhomogeneous opaque dielectric materials such as plastics. The selective absorption of certain wavelengths by colour pigments in the body of the material gives rise to a coloured body reflection. Such a surface typically has few pigment particles on the surface and hence the surface reflection component is unattenuated. Figure 3.8 shows the processes assumed by the dichromatic model. A thorough analysis of the dichromatic model and its various uses in computer vision is given by Klinker [66]. An obvious failing of the dichromatic model is the inability to model conductive materials which do not exhibit body reflection and attenuate the wavelength of interface reflections. Consideration of such materials, as well textiles, within the context of the dichromatic model has been attempted by Tominaga [130].

A number of studies have looked at modelling surface reflectance of particular materials. These include: vegetation [61]; snow and ice [140]; and bare soil [109]. Whilst these represent surface types that one would expect to encounter in natural scenes they have not been considered here. This work looks at the use of a single model that may be applicable to a wide variety of surface types.

Care should be taken when attempting to construct a hybrid reflection model based upon lobe models proposed by different authors. When considering a physical model of light reflection, it is imperative that the resulting total BRDF function has proper normalisation. A reflectance model that does not yield a correct energy balance is useless when considering the physical process of light reflection [141]. Failure to ensure this could cause the model to predict more light energy to be reflected than is incident on a surface. For this reason the reflectance model proposed by Ward is treated in isolation, though still with reference to the lobe model above. This model has been formulated to ensure proper conservation of energy and should not be used in conjunction with the other lobe models without proper normalisation.

3.4.1 Ideal Specular Reflection

A perfectly specular (mirror-like) reflecting surface reflects incident radiance only in the mirror direction defined by the law of reflection. Such a surface can be described in terms of a BRDF involving Dirac delta functions [12],

$$f_{spec} = \frac{\delta(\cos \theta_i - \cos \theta_r)}{\cos \theta_i} \delta(\phi_i - (\phi_r \pm \pi)), \quad (3.15)$$

where the Dirac delta functions, δ , have the following properties:

$$\begin{aligned} \delta(x) &= 0 \quad \text{for } x \neq 0, \\ \int \delta(x) dx &= 1, \quad \text{and} \\ \int f(x) \delta(x - a) dx &= f(a). \end{aligned} \quad (3.16)$$

Ideal specular surfaces are rarely encountered in natural scenes. Almost all surfaces contain surface imperfections which render the above consideration of specular reflection impractical. For this reason many reflection models neglect this portion of the total reflectance.

3.4.2 The Normal Lobe

Incident light reflected about the surface normal, independent of the incident direction, is contained within the *normal lobe*. Real-world surfaces can reflect a high proportion of incident energy into this lobe. The earliest and simplest normal lobe reflection model is that proposed by Lambert (1760) [71] and has since become known as *Lambert's law*. Stated in terms of a BRDF this gives

$$f_{norm}(\vec{\omega}_i \rightarrow \vec{\omega}_r) = \frac{1}{\pi}. \quad (3.17)$$

Lambert's model has been shown to approximate experimental data for a large set of materials. A complete explanation of Lambert's law has not yet been proposed [127]. It is generally thought that incident light undergoes repeated scattering within the bulk of the material (see Figure 3.8) [150]. This scattering causes the emergent radiation to be uniform about the surface normal (see [127] for a summary of the various normal lobe theories).

Studies by Oren and Nayer [101, 102] have shown that the Lambertian model is a poor approximation for rough surfaces. The Oren-Nayer reflectance model uses the Gaussian facet

slope distribution of Section 3.3.2 with each Lambertian facet having the same albedo. The model accounts for facet shadowing, foreshortening and interreflection. The Oren-Nayer model provides a useful alternative to the Lambertian model in the case of rough surfaces and reduces to the Lambertian model for smooth surfaces. The model has been shown to fit experimental data when combined with the Cook-Torrance model for the foreshatter lobe (see Section 3.4.3) [16]. Care must be taken, however, in combining this model with the backscatter models of Section 3.4.4, since an element of reflection within the backscatter lobe is taken into account by this model.

3.4.3 The Foreshatter Lobe

An early attempt to model the observed glossy reflections associated with the foreshatter lobe was made by Phong [108]. This was an empirical model which cannot be adequately expressed in terms of a BRDF, as presented in this chapter due to its failure to be bidirectional. The model assumes maximum reflectance along the perfect mirror direction, $\theta_r = \theta_i$. Reflectance falloff is approximated by $\cos^n \alpha$ where n is the surface's *specular-reflection exponent* and α is the angular difference between the considered reflection direction and the ideal mirror direction. Phong's model fails to account for Fresnel effects and assumes smooth surfaces. For this reason images generated using this model appear plastic and, in many cases, unnatural.

Torrance and Sparrow [134] used geometrical optics methods to consider the phenomenon of observed off-specular peaks and incident light attenuation by metals. They concluded that this was a result of roughened surfaces and Fresnel effects. They employed the Gaussian slope model of Section 3.3.2 and added a geometric attenuation factor to account for facet masking and shadowing. Cook and Torrance [14] adapted the model and applied it for use in computer graphics. They replaced the Gaussian slope distribution with the one proposed by Beckmann (see Section 3.3.2). The resulting foreshatter model accounts for many of the observed shading effects seen especially from conductive materials with surface finishes,

$$f_{fsc}(\theta_i, \phi_i; \theta_r, \phi_r; \lambda) = \frac{DGF(\lambda)}{\pi \cos \theta_i \cos \theta_r}, \quad (3.18)$$

where D is the Beckmann microfacet distribution of Section 3.3.2, F is the Fresnel term (see Section 3.2) and G is a geometric attenuation factor to account for surface shadowing and masking (see Section 3.3.3).

There are many alternative models for the foreshatter lobe. Hall provides a good overview

of the various models that may be used [36]. Despite these, the Cook-Torrance model has remained the most popular geometric model. Physical models (such as [4, 42, 41]) derived from a consideration of the wave nature of light and the molecular structure of the material produce good results but are complex. It is important that a model should adequately describe the process of reflection but not at the expense of undue complexity. In the context of machine vision, the time taken to evaluate a model for a single surface patch is of importance when considering algorithms that are required to apply repeated evaluations over a whole image, or a sequence of images.

3.4.4 The Backscatter Lobe

Observed natural shading of many materials has provided evidence for the backscatter lobe [22, 87]. A well studied example is that of the surface of the Moon and other planetary bodies whose reflectance peaks in the direction of incidence, sometimes termed the *opposition effect* [38, 121]. This effect is most apparent from particulate surfaces such as sand or dry soil. Experimental evidence suggesting the existence of the backscatter lobe can also be found in the data supplied by Torrance and Sparrow [134]. Models for the backscatter lobe are presented by [127] and [102]. They have not been considered in this work owing to their complexity and relative immaturity.

3.4.5 Ward's Reflection Model

The reflection model proposed by Ward [143, 141] has been derived empirically from reflectance data obtained using an imaging gonireflectometer. Despite not having been formulated from a thorough consideration of the physics of light reflection, Ward's model has been shown to provide an accurate simulation of light reflection in complex environments [35, 81].

Ward's model dispenses with the specific geometric attenuation and Fresnel terms adopted by many models in favour of a single normalisation factor. This ensures that the reflection function integrates predictably over the hemisphere of incident and reflection directions, within the limits defined by the conservation of energy. This is an important point to consider. Some reflection models, though derived using a physically based approach, fail to account for this and hence may provide inaccurate results when global illumination is considered.

In a similar approach to the model proposed by Cook and Torrance (see Section 3.4.3),

Ward uses a Gaussian process to model isotropic surface slope and hence glossy reflection into the forescatter lobe. This is combined with the Lambertian reflectance model covered in Section 3.4.2 to account for normal lobe reflection,

$$f_r(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{\rho_d}{\pi} + \rho_s \cdot \frac{1}{\sqrt{\cos \theta_i \cos \theta_r}} \cdot \frac{\exp(-\tan^2 \delta / \sigma_\alpha^2)}{4\pi\sigma_\alpha^2}, \quad (3.19)$$

where ρ_d is the diffuse reflectance into the normal lobe, ρ_s is the specular reflectance into the forescatter lobe, σ_α is the rms slope (see Section 3.3.2), δ is the *half angle*, and $\rho_d + \rho_s \leq 1$. The half angle is the angle between the surface normal and a vector bisecting the incident and reflection directions. Reflection into the backscatter lobe is not considered by this model. The inclusion of a specular term allows the modelling of perfectly smooth, glossy surfaces which do not scatter light.

Ward suggests that ρ_d values of less than 0.5 be used to model plastics, whilst ρ_s values greater than 0.5 be used for metallic surfaces. By making each of the reflectance terms a function of wavelength, the model can be extended to characterise reflection from coloured surfaces. In the case of plastics, coloured body reflection can be achieved by making ρ_d a function of wavelength, specular interface reflection being the colour of incident light, as described by the dichromatic model. For metals, both ρ_d and ρ_s are made a function of wavelength, hence providing the coloured specular reflection characteristic of metal surfaces. This may be performed provided $\rho_d + \rho_s \leq 1$. Experimental data also suggests that a practical range for σ_α is (0.0 – 0.2) [141]. The model can be further extended to model Fresnel effects by making the ρ_s term a function of both wavelength and geometry, though this is not considered here.

This work considers three variants of the basic reflection model given by Equation 3.19. The first requires that $\rho_d + \rho_s = 1$ and therefore has two free parameters: ρ_s and σ_α . The second variant provides an achromatic reflection model which allows the total reflectance of the surface to be given. This model therefore has three parameters: R , ρ_s , and σ_α where R is the total reflectance and $\rho_d = 1 - \rho_s$;

$$f_r(\theta_i, \phi_i; \theta_r, \phi_r) = R \left[\frac{\rho_d}{\pi} + \rho_s \cdot \frac{1}{\sqrt{\cos \theta_i \cos \theta_r}} \cdot \frac{\exp(-\tan^2 \delta / \sigma_\alpha^2)}{4\pi\sigma_\alpha^2} \right]. \quad (3.20)$$

Finally a colour reflection model is considered. This replaces the total reflectance R with a material colour vector \mathbf{C} to denote the reflectance of the surface at three discrete wavelengths. Again, $\rho_d + \rho_s = 1$ but how the colour vector is used in the model depends upon the value of ρ_s . If $\rho_s < 0.5$ the material is considered plastic and as such the colour of the glossy highlight is not affected by \mathbf{C} . If $\rho_s > 0.5$ the material is considered to be metallic and both ρ_s and ρ_d are made

functions of \mathbf{C} to provide coloured body reflection and coloured glossy highlights. Note that the colour vector \mathbf{C} is not normalised.

The reflectance function given in Equation 3.19 has been derived using reflection data obtained from a broad variety of material types. As such, it has not been designed to model reflection from a particular class of materials. This is a particular failing of the previous models presented in this chapter. Though alternative models provide accurate simulation of light reflection from surfaces they do not, in general, extend to cover a wide variety of surface types. Ward's model has been shown to characterise reflection from surfaces such as wood, metal, cardboard, plastic, ceramic and paper. One would expect to encounter such surface types in a natural scene.

Since Ward's model has been derived from the fitting of observed reflection data, it would appear suitable for the machine vision task of surface model estimation presented here. Reflection from the surface types mentioned above have been shown to fit the model in the laboratory situation and it would therefore be expected that similar results could be obtained from surfaces observed under natural illumination.

This model may be extended to characterise anisotropic surface reflection by the inclusion of further roughness parameters. Since such reflection is not considered here, the full anisotropic model is not presented.

3.5 Summary

This chapter has shown the development of models to describe surface reflection. These have been derived by consideration of the interaction of light with ideal surfaces. Practical models extend this theory to characterise reflection from surfaces which deviate from the ideal. As such, the models presented are able to account for scattering by material structures and by surface irregularities. These enable the modelling of some of the observed reflection phenomena which may be observed from many surfaces. It has also been shown that specific models exist to describe light reflection from a number of natural surfaces.

It has been shown that the reflection model due to Ward is most suited to the task presented here. It is able to model the reflection from a wide variety of surfaces and is computationally convenient. Since this model has been developed using an imaging gonireflectometer in laboratory

conditions it would appear to be suitable for the similar, image-based, system observing natural scenes presented here. Whilst the methods developed in this thesis do not pertain to any particular reflection model, the Ward model has been selected for experimental purposes.

Chapter 4

Natural Illumination

The appearance of an object is determined not only by the material properties discussed in Chapter 3 but also by the illumination to which the object is exposed. For example, a mirror viewed under a diffuse, uniform light source will appear very different to one observed under a single point source. It is also necessary to consider the spectral nature of the illumination. A material which absorbs radiation in the red region of the spectrum will appear to have no reflective properties when viewed under red illumination. In the case of natural illumination, an object illuminated by a cloudy, overcast sky may well look different to an identical object viewed under a clear blue sky with an unobstructed sun. This chapter is concerned with the effects of natural illumination on objects and how their subsequent appearance may be modelled.

Natural illumination, or more specifically daylight illumination, can be considered as being composed of two components: direct solar and diffuse sky light illumination. The magnitude and distribution of each of these components is determined by factors such as location, time and weather conditions. In order to predict how an object will appear given these parameters, a reliable and accurate daylight illumination model is required. Such models have been used extensively in the energy and architecture communities and their use is now becoming more common in computer graphics and machine vision.

4.1 Direct Solar Illumination

4.1.1 Solar Geometry

The earth rotates about the sun approximately once every $365\frac{1}{4}$ days in an almost circular path. The sun is located slightly away from the centre of the earth's orbit which causes the solar intensity to be about 7 per cent higher in January than July. The earth also spins about its axis every 24 hours giving diurnal variation in solar intensity. The earth's axis of rotation is tilted by 23.5° relative to its plane of motion and this causes seasonal variation in sun position. Therefore, the position of the sun in the sky hemisphere, and hence solar intensity, is determined by date, time and global location. The average solid angle subtended by the sun at a point on the earth is 6×10^{-5} steradians [103].

At any time, the position of the sun, S , in the sky hemisphere for a global location can be specified by spherical polar co-ordinates (θ_s, ϕ_s) , where θ_s is the solar zenith and ϕ_s is the solar azimuth. These angles represent the location of the sun relative to a vertical direction and the north direction at a solar hour angle, h . The solar hour angle is calculated from the local solar time (LST). LST accounts for location longitude and uses the *equation of time* to compensate for the eccentricity of the earth's orbit [112]. Local solar time is such that at 12:00 hours, the sun is due south of the location. Given a local solar time, the hour angle h is given by,

$$h = (LST - 12) * 15^\circ. \quad (4.1)$$

Figure 4.1 shows the horizontal plane at a global location O with latitude L° in the northern hemisphere. On an equinox (April 21 or September 21) the declination of the earth is such that the sun appears to move within a plane tilted by L° along the west-east axis. Such a plane is termed the *equatorial plane* for the location O . On other dates the declination of the earth, relative to the sun, is such that the plane of apparent motion of the sun is shifted by an angle equal to the declination d . This plane of apparent motion is termed the *declination plane* for the location. The direction V is an *up* direction which is perpendicular to the horizontal plane at the location O . The direction P is a direction perpendicular to the sun's plane of apparent motion. Note that the possible declination planes for a location are parallel to the equatorial plane.

Looking at the spherical triangle PVS , which is composed of great circles, the solar zenith,




Figure 4.1: Solar geometry (follows Rapp [112]).

θ_s , is given by the law of cosines,

$$\begin{aligned}\cos \theta_s &= \cos(90 - d) \cos(90 - L) + \sin(90 - d) \sin(90 - L) \cos h \\ &= \sin d \sin L + \cos d \cos L \cos h\end{aligned}\quad (4.2)$$

where

$$\sin d = -\cos \left[(D_s - 1) \frac{180^\circ}{182.6} \right] \sin(23.45^\circ) \quad (4.3)$$

where $D_s = 1$ on December 21, and $D_s = 365$ on December 20.

The solar azimuth angle, ϕ_s , is defined as the angle between the north direction and the projection of the sun onto the horizontal plane. Applying the cosine law to the spherical triangle *PVS* the azimuth angle can be found as,

$$\begin{aligned}\cos(90 - d) &= \cos(90 - L) \cos \theta_s + \sin(90 - L) \sin \theta_s \cos \phi_s \\ \cos \phi_s &= \frac{\sin d - \sin L \cos \theta_s}{\cos L \sin \theta_s}\end{aligned}\quad (4.4)$$

where $\sin d$ is given by Equation 4.3.

4.1.2 Solar Intensity

Owing to the elliptical trajectory of the earth, the distance between the sun and earth changes throughout the year. The extraterrestrial irradiance due to the sun on a surface normal to the sun, when the earth is at a mean distance from the sun, is termed the *solar constant*. Traditionally this has been taken to be 1353 W m^{-2} [17]. Measurements suggest, however, that a value of 1377 W m^{-2} is more appropriate [112]. The extraterrestrial solar intensity, N_{ext} , for a given day in the year may be estimated from the solar constant [112],

$$N_{ext} = \frac{S(1 + \epsilon \cos \phi)^2}{(1 - \epsilon^2)} \quad (4.5)$$

where, S is the solar constant, ϵ is the eccentricity of the earth's orbit ($=0.01672$) and ϕ is given by,

$$\phi = (D_j - 2) \frac{360^\circ}{365.2} \quad (4.6)$$

where D_j is the day of the Julian year. Figure 4.2 shows a graph of extraterrestrial solar irradiance on a surface normal to the sun's rays on the twenty first day of each month as given by Equation 4.5, assuming a solar constant of 1377 W m^{-2} .


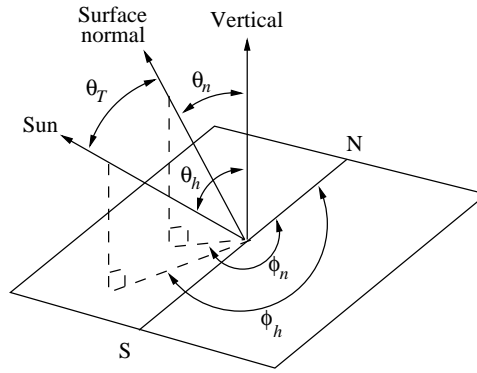


Figure 4.2: Extraterrestrial solar irradiance incident on a surface normal to the direction of the sun's rays on the twenty first day of each month, calculated using Equation 4.5.

The intensity of the sun at the earth's surface is reduced from the extraterrestrial amount due to atmospheric scattering and absorption. Section 4.2 discusses the scattering processes which solar radiation undergoes in order to form the diffuse sky light. The amount of direct solar radiation lost to scattering and absorption is dependent upon the path taken by the radiation through the atmosphere. The level of radiation received by a surface outside of the atmosphere is referred to as *air mass zero*. When the sun is directly above a surface at sea level, the radiation received is *air mass one*. When the sun is located such that the radiation passes through a greater amount of atmosphere, the effective air mass is increased. The effective path length to sea level through the atmosphere, l , and hence air mass, at any time is given by $l = \sec \theta_h$, where θ_h is the solar zenith. The fraction of light absorbed by the atmosphere over a path length of l is given by $\exp(-\beta(\lambda)l)$, where $\beta(\lambda)$ is the *extinction coefficient* for radiation of wavelength λ . The extinction coefficient for radiation of all visible wavelengths may be approximated by a single value. A reasonable approximation for a clear sky is $\beta(\lambda) = 0.431$ [112]. Increased scattering due to moisture or pollution will increase the extinction coefficient such that direct solar irradiance becomes negligible as the atmosphere becomes more overcast.

The direct solar irradiance received by a plane at the earth's surface, oriented such that it is normal to the direction of the sun's rays is,

$$N = N_{ext} \exp(-\beta(\lambda) \sec \theta_h). \quad (4.7)$$

Figure 4.3: *Tilted surface geometry.*

A tilted plane will receive more or less irradiance direct from the sun. Maximum irradiance will be received when the surface is tilted such that the sun is in line with the surface normal. Figure 4.3 shows the geometry under consideration for a tilted surface illuminated by direct solar irradiance. The cosine of the angle between the sun direction and the tilted surface normal θ_T is given by,

$$\cos \theta_T = \cos \theta_n \cos \theta_h + \sin \theta_n \sin \theta_h \cos(\phi_h - \phi_n). \quad (4.8)$$

Direct solar irradiance incident on the tilted plane, D_T , is therefore,

$$D_T = N \cos \theta_T. \quad (4.9)$$

Figure 4.4 shows the direct solar irradiance on a horizontal surface and a surface tilted 30° to the south for a location with latitude 31°N on a clear day. Values were calculated using Equations 4.7 and 4.9 assuming noon air mass of 1.7, constant extinction coefficient of 0.23, and solar constant of 1377 W m^{-2} . Measured data was recorded at a location with the same latitude, on a clear day in 1977, using a normal incidence pyrheliometer [113].

4.1.3 Spectral Distribution

The sun emits radiation over a wide range of wavelengths. Of particular importance is radiation within the visible spectrum between $0.4\mu\text{m}$ and $0.7\mu\text{m}$. Extraterrestrial data suggests that the radiation emitted by the sun compares closely with the expected radiation from a black body at 5762 K [112]. Scattering and absorption by the atmosphere of particular wavelengths causes the terrestrial spectral distribution of direct solar energy to be shifted (see Section 4.2). The spectral




Figure 4.4: Measured and predicted solar irradiance on a horizontal and tilted surface for a location with latitude 31°N on January 4 1977. The tilted surface is inclined 30° to the south. Predicted solar irradiance is calculated using Equations 4.7 and 4.9 assuming noon air mass of 1.7, constant extinction coefficient of 0.23, and solar constant of 1377 W m^{-2} [113].

composition of terrestrial sunlight depends upon the level of scattering and hence upon the composition of the atmosphere. Measured data shows that the terrestrial solar spectrum correlates to black body radiation at 4000 K [46, 48, 56, 128]. As the sun's altitude increases the spectrum can be compared to a correlated colour temperature of 5500 K [48].

4.2 Diffuse Sky Light Illumination

To an observer on the ground, sky light appears to emanate from a hemisphere of directions. The diffuse nature of sky light is due to atmospheric scattering of radiation incident from the sun. Scattering occurs when light strikes particles suspended in a medium and can be due to reflection or refraction on meeting the particles. Two widely used models account for scattering by particles of varying sizes [85]:

Rayleigh Scattering: This scattering model is applicable for particles which are smaller than the wavelength of light under consideration. This model is of particular interest when considering scattering by air molecules.

Mie Scattering: Scattering by particles which are larger than the wavelength of light is described by Mie scattering theory. This model accounts for scattering by aerosols. This model is applicable when considering moisture, dust and pollutants suspended in the atmosphere.

The extent to which light is scattered in each of the above cases is dependent upon the wavelength of light being considered. The characteristic blue sky observed on clear days is due to Rayleigh scattering, and the absence of Mie scattering by moisture particles. The colour of the sky shifts towards longer wavelengths as the level of moisture increases. Pollutants also play an important role in determining the overall scattering of solar radiation by the atmosphere.

Computer graphics and machine vision systems have generally used a constant ambient light to model the diffuse nature of sky light. This is a poor approximation because sky light, especially clear sky light, has high spatial variation. The following section provides a more detailed description of sky light which enables a more accurate model to be developed.




Figure 4.5: Sky luminance geometry.

4.2.1 Sky Light Distribution Models

The Commission Internationale de L'Éclairage (CIE) has formulated standard luminance distribution models for various reference skies. These were originally commissioned in response to a need by the illumination engineering community for a set of standard reference skies for lighting calculations. Two such models have been formulated, namely the *CIE standard clear sky* and *CIE standard overcast sky*. These two reference skies represent luminance distributions for extreme ideal weather conditions. To extend these models to account for skies between these two extremes, a general or intermediate sky needs to be defined. The CIE have yet to approve a general sky model which is applicable to all global locations and atmospheric conditions but various models have been proposed [106, 107].

Figure 4.5 shows the sky hemisphere geometry required for calculating the relative luminance of a discrete sky point. The solar zenith, θ_s , and azimuth, ϕ_s , can be determined using the equations given in Section 4.1.1.

CIE Standard Clear sky

The clear sky luminance, L_{cl} , of a point in the sky hemisphere with zenith angle θ_p and azimuth angle ϕ_p is,

$$\frac{L_{cl}(\theta_s, \theta_p, \zeta)}{L_{zcl}(\theta_s)} = \frac{\gamma(\theta_p)f(\zeta)}{\gamma(0)f(\theta_s)} \quad (4.10)$$

where

$$L_{zcl}(\theta_s) = \text{clear sky zenith luminance}, \quad (4.11)$$

$$\gamma(\theta) = 1 - e^{(-0.32/\cos\theta)}, \quad (4.12)$$

$$f(\zeta) = 0.91 + 10e^{-3\zeta} + 0.45\cos^2\zeta, \quad (4.13)$$

$$\cos(\zeta) = \cos\theta_s \cos\theta_p + \sin\theta_s \sin\theta_p \cos|\phi_s - \phi_p|. \quad (4.14)$$

The clear sky model can be further extended to account for atmospheres polluted with particulates. Such extensions typically use a sky turbidity factor to account for the level of atmospheric pollution [18]. Such models are particularly suitable for urban or industrial regions. Regional models also exist to account for the sky luminance distribution for a particular location, for example [90].

Figure 4.6 shows sky point luminance relative to the sky zenith luminance for clear skies with solar altitudes of 22.5° , 45° and 67.5° . Each greyscale image is accompanied by a false colour image showing the pattern of relative luminance over the whole sky hemisphere.

A general equation for the absolute clear sky zenith luminance has yet to be standardised by the CIE. A popular model, which is applicable to a variety of global locations, suitable for $\theta_s > 30^\circ$, is given by Kittler [64],

$$L_{zcl}(\theta_s) = 300 + 3000 \cot\theta_s \quad [\text{cd m}^{-2}]. \quad (4.15)$$

Figure 4.7 shows the overcast sky zenith luminance as predicted by this model. As with the luminance distribution model, a variety of alternative zenith models exist to account for differing conditions. In particular, models attempt to account for pollution and high turbidity. An overview of possible alternative models is given by [18].

CIE Standard Overcast Sky

The overcast sky luminance, L_{oc} , of a point in the sky hemisphere with zenith θ_p and azimuth ϕ_p is given by the CIE standard overcast sky model,

$$\frac{L_{oc}(\theta_p)}{L_{zoc}(\theta_s)} = \frac{1 + 2\cos\theta_p}{3} \quad (4.16)$$




Figure 4.6: Greyscale images of clear sky luminance relative to zenith as modelled by Equation 4.10 for: (a) $\theta_s = 67.5^\circ$; (b) $\theta_s = 45^\circ$; (c) $\theta_s = 22.5^\circ$; and false colour images for: (d) $\theta_s = 67.5^\circ$; (e) $\theta_s = 45^\circ$; (f) $\theta_s = 22.5^\circ$.

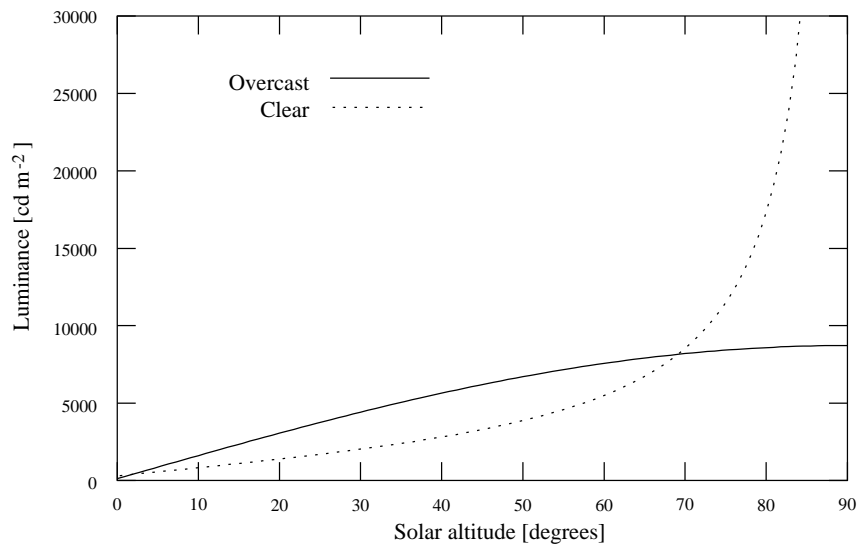


Figure 4.7: Sky zenith luminance for clear and overcast skies predicted by Equation 4.15 and 4.17.

where $L_{zoc}(\theta_s)$ is the overcast sky zenith luminance. Note that, unlike the clear sky model, the luminance of a sky point does not depend on its position relative to the sun. Also note that the distribution of relative illuminance is independent of the solar zenith angle.

Figure 4.8 shows sky point luminance relative to the sky zenith luminance for an overcast sky. The greyscale image is accompanied by a false colour image showing the pattern of relative luminance over the whole sky hemisphere.

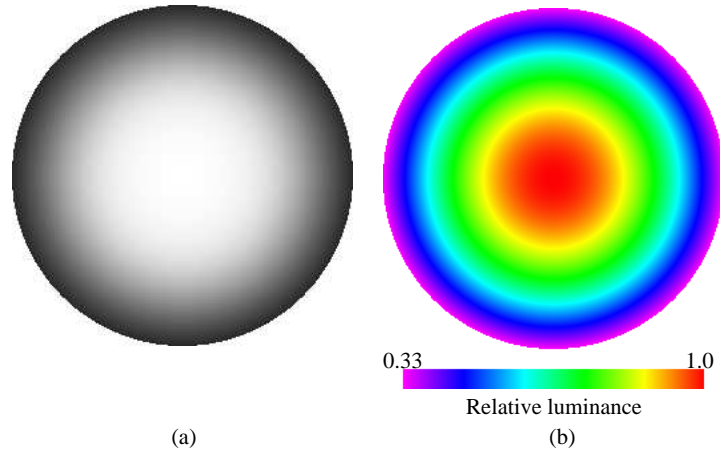


Figure 4.8: (a) Greyscale overcast sky luminance relative to zenith for sun zenith angle of 45° as modelled by Equation 4.16, and (b) false colour image of the same sky.

As with the clear sky model, there is no standard formula for calculating the overcast sky zenith luminance. A model which has been shown to match measured data is given by Krochmann and Seidl, [70]:

$$L_{zoc}(\theta_s) = 123 + 8600 \cos \theta_s \quad [\text{cd m}^{-2}]. \quad (4.17)$$

Figure 4.7 shows the overcast sky zenith luminance as predicted by this model.

Intermediate Sky

Models for describing intermediate skies have yet to be standardised by the CIE. Such a sky model would need to account for the large variety of weather conditions that are possible between the ideal clear and overcast skies presented above. A variety of methods for describing such intermediate skies are discussed in the CIE daylight report [18]. The method used in this work is the simple combination proposed by Gillette and Kusuda [29]. This model uses a *sun probability*,

ϕ , to interpolate between the CIE clear and CIE overcast sky models,

$$L_{in} = (1 - \phi)L_{oc} + \phi L_{cl}, \quad (4.18)$$

where L_{oc} and L_{cl} are the CIE luminance distributions for overcast and clear skies respectively. A clear sky has a high sun probability, hence $\phi = 1$. Similarly, an overcast sky is modelled using $\phi = 0$. This model assumes that cloud cover is homogeneous. Models to account for inhomogeneous skies, with patches of cloud and blue sky, have been considered but are not sufficient for practical use at present [135].

Figure 4.9 shows sky luminance relative to the sky zenith luminance for intermediate skies with a solar altitude of 45° , calculated using Equation 4.18. The skies shown have sun probabilities of 0.25, 0.5 and 0.75 and each greyscale image is accompanied by a false colour image showing the pattern of relative luminance over the whole sky hemisphere.




Figure 4.9: Greyscale images of intermediate sky luminance relative to zenith for sun zenith angle of 45° as modelled by Equation 4.18 for: (a) $\phi = 0.25$; (b) $\phi = 0.5$; (c) $\phi = 0.75$; and false colour images for: (d) $\phi = 0.25$; (e) $\phi = 0.5$; (f) $\phi = 0.75$.

This intermediate sky model has been evaluated by Littlefair, who compared the model to illuminance data recorded near London between July 1991 and January 1993 [77]. The results suggest that, over the whole sky hemisphere, the model has an RMS error of 43 per cent when compared to measured data for skies over this period. The model performs well when considering cloudy and intermediate skies but performs relatively poorly when considering clear skies. Littlefair shows that, in the case of intermediate skies ($\phi = 0.5$), the intermediate sky model predicts increased sky luminance values. This may be due to the assumption of homogeneous cloud cover. Despite these deficiencies, the model is simple to compute and requires few parameters. Determining a sun probability value for an observed sky is simple when compared to calculating the parameters that are required for more accurate intermediate sky descriptions.

4.2.2 Sky Light Luminous Efficacy

The above models provide photometric luminance values for discrete sky points. In order to convert these to radiometric quantities the luminous efficacy of the light must be considered. The luminous efficacy relates the luminous intensity of the sky to a radiant intensity. The luminous efficacy of light is dependent upon wavelength since it is based upon the photopic response of the human visual system (see Appendix A).

Assuming a luminous efficacy which is constant with respect to wavelength, data suggests that a value of 150 lm W^{-1} for clear skies and $115 - 125 \text{ lm W}^{-1}$ for overcast skies is suitable [48, 70]. The change in luminous efficacy between clear and overcast skies is due to the shift in spectral distribution and hence the perceived brightness.

By applying the above efficacy values to the appropriate zenith luminance models for each reference sky, the zenith radiance may be determined. When combined with the relative distribution models, the radiance of any point in the sky hemisphere may be determined.

4.2.3 Spectral Distribution

The scattering processes described in Section 4.2 show that the specular distribution of sky light is dependent upon the extent and type of scattering that incident solar radiation undergoes. Klassen [65] and Inakage [55] have shown that, by modelling the various scattering processes, the spectral

attenuation of light due to the atmosphere can be modelled. A similar approach has been used by Tadamura et al. [125] which they show to match closely the CIE illuminance distributions. An alternative, and less complex, approach is to use measured data, which suggests a correlated colour temperature of 4500 – 7000 K for overcast skies and 10^4 – 10^5 K for clear skies [46, 48].

4.2.4 Surface Irradiance

Each of the sky models described above provides the radiance incident from a differential solid angle. To calculate the total irradiance incident on a surface from the sky it is necessary to integrate the distribution functions over the visible sky hemisphere. In the case of a horizontal surface element this requires integration over the whole sky hemisphere. In the case of a tilted surface it is necessary to determine the regions of the sky which are visible above the horizon. In scenes where there may be shadowing, it is necessary to calculate which regions of the sky are masked by other objects before computing the sky irradiance. As shown in Appendix A, surface irradiance is calculated by integrating incident radiance over the hemisphere of visible directions [32]. In the case of a surface illuminated by sky light this integration becomes,

$$E = \int_0^{2\pi} \int_0^{\frac{\pi}{2}} L(\theta, \phi) \cos \theta d\theta d\phi \quad (4.19)$$

where $L(\theta, \phi)$ is the sky radiance determined using the models described in Section 4.2. To account for tilted surfaces and shadowing a sky visibility term would need to be added into Equation 4.19 for each differential direction.

The integration of available light incident on a surface is a common problem in computer graphics and a variety of solutions have been proposed. Numerical integration of the irradiance is desirable but comes at the expense of computational time. Equation 4.19 can be approximated using methods derived from computer graphics and heat transfer studies. Three such methods are considered here.

Hall's Hemispherical Integrator

Hall has presented a simple method for calculating the irradiance incident on a surface from a hemisphere of directions [36]. This technique is easily adapted to the computation of sky irradiance for tilted and shadowed surfaces. The surface element, for which the calculation is to be

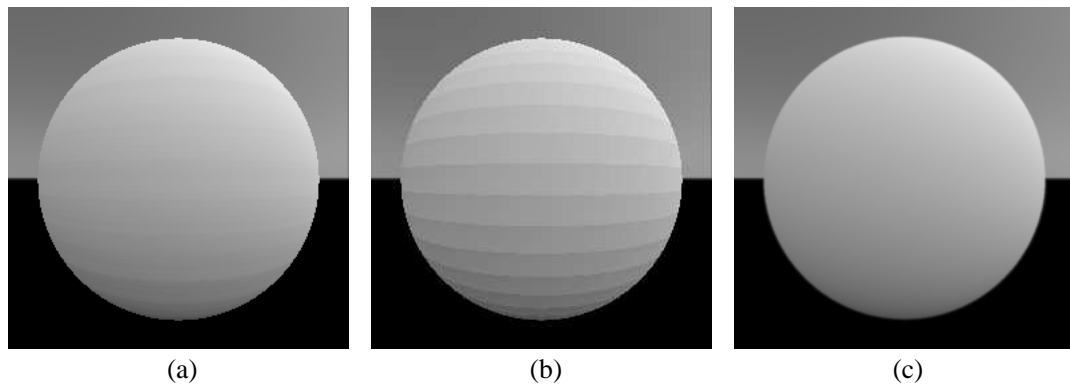


Figure 4.10: Irradiance from a CIE clear sky calculated on the surface of a sphere. The camera is viewing the sphere from the north and the sun is located due west with a zenith angle of 30° : (a) Irradiance calculated using Hall's hemispherical integrator with an angle increment of $\pi/16$ (11.3 minutes); (b) calculated using Nishita's band method using 24 bands (25.6 minutes); (c) calculated using Ward's Radiance lighting simulation system using high image quality parameters (3.5 minutes).

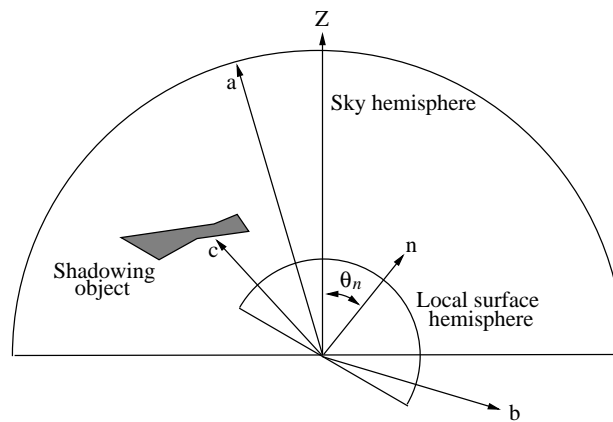


Figure 4.11: Use of a local hemisphere of incident directions to sample the sky hemisphere.

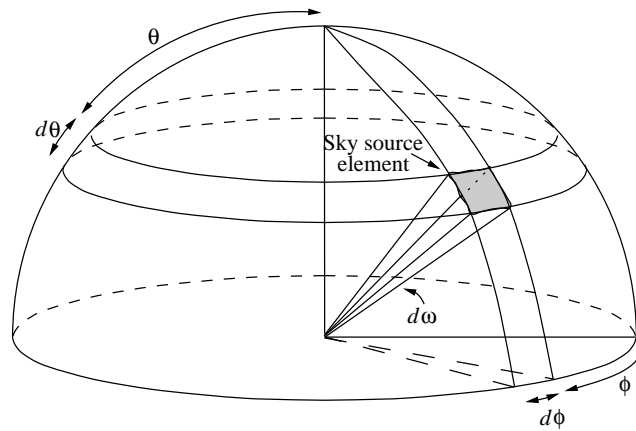


Figure 4.12: Geometry for sky irradiance calculation using Hall's hemispherical integrator [36].

performed, is assumed to be in the centre of the base of the sky hemisphere. This is a fair assumption assuming the size of the sky hemisphere to be large when compared to the size of the element. A local hemisphere of visible directions is placed over the surface patch under consideration. This local hemisphere is positioned such that its zenith is coincident with the surface normal. The local hemisphere of directions is sampled and these directions are mapped to the sky hemisphere to provide radiance values. Figure 4.11 shows the use of a local hemisphere for sampling the global sky hemisphere. The figure shows a sample direction *a* which successfully samples the sky hemisphere and hence contributes irradiance. Sample direction *b* is below the horizon and hence irradiance from the sky does not contribute from this direction. Sample direction *c* is incident on another scene object and is therefore shadowed from sky irradiance in that direction.

To calculate irradiance the local surface hemisphere is divided into discrete elements such that the angle increments in the longitudinal and latitudinal directions are equal ($d\theta = d\phi$) (see Figure 4.12). The centre of the hemisphere element is sampled to determine the possible sky radiance of the element. When the direction is below the horizon or is incident on another object in the scene, the sky element is assigned a radiance of zero. Equation 4.19 is now approximated by summing the contribution from each hemisphere element

$$E \approx \sum_{i=0}^n \sum_{j=0}^{4n} L(i\pi/2n, j\pi/2n) \cos(i\pi/2n) d\omega \quad (4.20)$$

where $d\omega$ is the projected solid angle of the element source and n is the number of divisions made.

Figure 4.10(a) shows the irradiance incident on a sphere calculated using this sampling method. Slight spatial aliasing can be seen, owing to insufficient sampling of the hemisphere. The

technique does, however, provide a good approximation of the irradiance incident on the surface.

Nishita's Band Integrator




Figure 4.13: Geometry for sky irradiance calculation using Nishita's band source integration method [96].

A popular computer graphics technique for calculating sky light irradiance is due to Nishita and Nakamae [96]. This method aims to reduce the time spent in determining which regions of the sky are visible to an element. As with Hall's method above, the hemisphere of possible incident directions is placed above the surface element. This local hemisphere is divided into bands which are treated as transversely uniform band luminaires (see Figure 4.13). When determining sky visibility, only the mid-line of each band is sampled. The visible band region is then integrated to determine the irradiance received by the element from the sky band source. The irradiance received from band l is given by

$$E_l = d_l \int_0^\alpha L(\alpha, \delta) \sin^2 \alpha d\alpha. \quad (4.21)$$

where $d_l = (\cos \delta'_l - \cos \delta'_{l+1})$ and $\delta'_l = \delta_l - \Delta\delta$. The total irradiance is calculated by summing the contribution from each of the band sources,

$$E \approx \sum_{l=0}^n E_l. \quad (4.22)$$

Figure 4.10(b) shows the irradiance incident on a sphere calculated using this method. Aliasing in the form of bands can be seen owing to insufficient sampling of the hemisphere. This is due to the technique being originally intended for planar convex polygons as opposed to the curved surface presented by a sphere. The approximation could be improved by dividing the local hemisphere into more bands but this would be at the expense of increased computation time,

though the method has since been modified so as to improve rendering time [19]. Results using this method are poor when compared to the simple method advocated by Hall.

Ward's *Radiance* Lighting Simulation System

An alternative to the deterministic sampling techniques described above is to randomly sample the local hemisphere of visible directions shown in Figure 4.11. Provided that enough samples are taken and the sample density is uniform, Equation 4.19, and hence irradiance incident on the surface, can be approximated. Such a stochastic method is often termed a Monte Carlo approach, due to the inherent random nature [32]. The use of Monte Carlo techniques to determine surface irradiance is ill-advised when considering daylight [120]. This is due to the high spatial variance of the illumination, especially when considering clear and intermediate skies. When direct irradiance from the sun is also considered, the use of stochastic methods is to be avoided owing to the small size of the sun relative to the sky hemisphere. In order to ensure that the high radiance regions of the hemisphere are not missed by the random sampling, many thousands of samples would have to be taken. Despite this, a lighting simulation system based on Monte Carlo methods, developed by Ward, has proved successful for the simulation of daylight illumination [142, 81].

The *Radiance* lighting simulation system [142] was originally designed for artificial interior lighting calculations and has been extended to account for exterior and daylight scenes. The system uses a hybrid deterministic and stochastic ray tracing technique to provide physically accurate lighting calculations for an extensive range of surfaces and illuminants. The system primarily uses a Monte Carlo technique to account for diffuse illumination such as sky light. Since direct solar and clear sky illumination have a high variance, deterministic methods are also used. In the case of a clear sky with a sun, Monte Carlo methods would be used to sample the majority of the sky hemisphere. Knowing that the regions of high variance in the sky hemisphere exist near the location of sun the majority of samples are made in this region. This hybrid technique allows for accurate calculations with the advantage of low cost. The number of rays cast using this method is considerably lower than in the methods presented previously. Mardaljevic [81] has shown a good correspondence between the results produced by the *Radiance* system and actual daylight measurements for a variety of sky conditions.

Figure 4.10(c) shows integrated irradiance on the surface of a sphere incident from a clear

sky without a sun. This method does not exhibit the aliasing effects which are produced by the Hall and Nishita methods. However, this improved calculation of surface irradiance does not come at the expense of computation time, as indicated by the comparative timings¹ shown in Figure 4.10. The time taken to compute the irradiance values is significantly improved over the two previous methods. Figure 4.14 shows a reconstruction of the daylight sequence shown in Figure 1.1 produced using the *Radiance* lighting simulation system.



Figure 4.14: *Reconstruction of the daylight image sequence shown in Figure 1.1, produced using the Radiance lighting simulation system.*

4.3 Summary

Natural illumination can be considered as being the combination of two light sources. These are direct irradiance from the sun and diffuse irradiance from the sky hemisphere. The relative contribution of each of these is determined by the atmospheric conditions which affect the level of direct solar illumination reaching the terrestrial surface. Each of these sources of illumination have been treated separately in the models presented.

This chapter has shown that models developed in the solar energy community may be used directly to provide a physically based model of natural illumination. These can be used in conjunction with computer graphics techniques to calculate the total irradiance on a surface, and hence reflected radiance. Such calculations allow for the interpretation of reflection measurements made by the methods developed in this thesis.

¹Timings obtained on SiliconGraphics Indy.

Chapter 5

A Sensor Model

Accurate measurement of the radiance reflected by a surface in the direction of a sensor can be achieved using a specialised calibrated device such as a gonioreflectometer [89]. Measurements obtained using such devices can be used directly for the development of a bidirectional reflectance-distribution function (BRDF) model. Machine vision has, however, relied on the use of CCD¹ cameras and digitised photographic images. Such image based sensors provide cost effective, compact and robust technology that can be used in a wide variety of situations. Before measurements are made from the images provided by such cameras, care must be taken to consider the processes by which the resultant image is formed. This chapter will focus upon the characteristics of CCD based devices in the context of measuring scene radiance. Similar consideration need be applied to digitised photographic images as these are typically scanned using a CCD device, and hence, similar sources of error may be encountered. In either case the limitations imposed by the camera optics are the same.

The sensor model developed here is based upon that of a CCD sensor device with appropriate optics to capture an image of a scene. Such a camera has inherent sources of error and these shall be discussed in this chapter. The significance to which these errors affect the resultant image can be determined by prior calibration of the camera. This chapter will discuss both radiometric and projective calibration. The sensor model developed here may be adapted to allow the analysis of data obtained using alternative image based sensors.

¹Charge-Coupled Device.

5.1 Projection of the Scene


To determine which regions of a scene contribute to an area of the image it is necessary to model the projection of the three dimensional scene onto the image plane [24]. Ideally, this perspective projection of the 3D scene onto the 2D image plane would be achieved using a pin-hole aperture. The use of such an ideal, infinitesimally small, aperture provides an undistorted projection of the scene onto the image plane, with all objects in sharp focus. This model of projection is that which is commonly used in the computer graphics community but the model only approximates the processes by which a scene is projected onto the image plane by a practical lens.

The sensitivity of a sensor determines the flux necessary to produce an image. The aperture may therefore need to be enlarged to allow sufficient light onto the image plane. This increase in aperture size compromises the pin-hole model. So that objects are brought to focus on the image plane, a lens is used in conjunction with the aperture. Only objects within the *principal plane* of the lens are brought to sharp focus, whilst those within the *depth of field* of the lens are brought to apparent focus on the image plane. Objects beyond the depth of field are represented by a region on the image plane, termed the *circle of confusion*. The perspective projection of objects within the depth of field can be approximated by the pin-hole model, the parameters of which may be determined by prior calibration (see Section 5.4).

To determine which regions of the scene contribute to an individual image pixel it is necessary to perform the inverse perspective projection. This can be achieved using the common computer graphics technique of tracing a ray backwards through the projection and into the scene [30, 31]. The surface patch which contributes to the selected pixel is determined by intersecting this ray with the scene. This technique requires *a priori* knowledge of the scene geometry within a defined co-ordinate system.

5.2 Image Plane Irradiance

A camera lens provides a system of optics through which light is focused onto the image plane. Figure 5.1 shows a single lens of focal length f_p and aperture of diameter D bringing an image of a surface patch, dA_o , to focus on the sensor's image plane. Image plane irradiance, E_p , is due to

Figure 5.1: *Imaging geometry for a simple camera model.*

the flux within the projected solid angle of the patch passing through the aperture of the lens,

$$E_p = \frac{d\Phi}{dA_p}. \quad (5.1)$$

This flux is due to the reflected radiance from the patch within the solid angle,

$$d\Phi = dA_o \int_{\Omega} L_r(\theta_r, \phi_r) \cos \theta_r d\omega, \quad (5.2)$$

where Ω is the projected solid angle subtended by the aperture. Substituting into Equation 5.1 gives:

$$E_p = \frac{dA_o}{dA_p} \int_{\Omega} L_r(\theta_r, \phi_r) \cos \theta_r d\omega. \quad (5.3)$$

Comparing the projected solid angles subtended by dA_o and dA_p at the aperture shows,

$$\frac{dA_o \cos \theta_r'}{f_o^2} = \frac{dA_p \cos \alpha}{f_p^2}. \quad (5.4)$$

Image plane irradiance can therefore be defined in terms of the reflected scene radiance and the lens system,

$$E_p = (f_o/f_p)^2 \cos \alpha \int_{\Omega} L_r(\theta_r, \phi_r) (\cos \theta_r / \cos \theta_r') d\omega. \quad (5.5)$$

Assuming the aperture of the lens to be small relative to the distance from the object allows the ratio of the cosines, $(\cos \theta_r / \cos \theta_r')$, to be unity. It can also be assumed that the reflected radiance from the surface patch will tend to be constant over the entire solid angle. These assumptions allow Equation 5.5 to be approximated [51]:

$$\begin{aligned} E_p &= \left(\frac{\pi}{4}\right) \left(\frac{D}{f_p}\right)^2 \cos^4 \alpha L_r(\theta_r, \phi_r) \\ &= l(\alpha) L_r(\theta_r, \phi_r). \end{aligned} \quad (5.6)$$

The reflected radiance, $L_r(\theta_r, \phi_r)$, from a surface is due to incident irradiance from light sources and interreflection from other surfaces being reflected in accordance with the surface's reflection characteristics. Chapter 3 has shown that a surface's reflection characteristics are most generally defined using a BRDF, f_r . The reflected radiance due to illumination incident from the hemisphere of possible directions about the surface normal is,

$$L_r(\theta_r, \phi_r) = \int_0^{2\pi} \int_0^{\frac{\pi}{2}} f_r(\theta_i, \phi_i; \theta_r, \phi_r) L_i(\theta_i, \phi_i) \cos \theta_i d\theta_i d\phi_i. \quad (5.7)$$

This double integral equation has a similar form to that of Equation 4.19 for the calculation of surface irradiance due to sun and sky light. Equation 5.7 may be solved using the same computational methods used to compute incident irradiance (see Section 4.2.4).

The image plane irradiance due to the illuminants and scene objects, given the simple camera model above, is therefore,

$$E_p = l(\alpha) \int_0^{2\pi} \int_0^{\frac{\pi}{2}} f_r(\theta_i, \phi_i; \theta_r, \phi_r) L_i(\theta_i, \phi_i) \cos \theta_i d\theta_i d\phi_i. \quad (5.8)$$

In the case of a scene illuminated by sun and sky light the incident radiance, L_i , can be provided by the illumination models presented in Chapter 4 and from surface interreflection. This work does not consider the contribution due to surface interreflection and hence illumination is due only to visible sun and sky light.

5.3 Production of the Digital Image

An image represents the spatial variation of radiance incident in the direction of the camera from the scene. For the purposes of machine vision, this continuously varying radiance is discretised to provide a numeric representation of the scene in the form of an array of pixel values. This section considers how such an image is formed in a CCD camera. The models presented in this section may equally be applied to cameras using photographic media, in which case the discretisation takes place when the image is scanned to provide pixel values.

A CCD chip, used in the majority of commercial imaging systems, is segmented into a grid of individual sensor sites. When a photon strikes a CCD site an electron is generated in the silicon structure. Each site will integrate photons over the duration that the site is exposed by the mechanical shutter or, in the case of an electronic shutter, until the photon-generated electrons are

collected from the site. The charge generated at each site is read off and amplified by the CCD chip circuitry to provide a signal suitable for processing. The resultant signal is quantised to provide a pixel value within a range of discrete values, the *dynamic range*. A thorough description of the processes by which a CCD chip generates a signal is provided by Healey and Kopardy [44].

An achromatic camera will provide a grey-scale image of the scene. The grey levels in the image are due to scene radiance integrated over the wavelengths of light to which the CCD sensor is responsive. To provide a colour representation of the scene the light is selectively filtered, either over the individual sensor sites or over different CCD chips, to provide a composite image. This section presents pixelised image formation for both achromatic and colour cameras.

5.3.1 Achromatic Image Formation

The number of photons striking an individual sensor site on a CCD chip is proportional to the incident flux over the duration to which the site is exposed to flux. The total flux incident upon an individual sensor site at grid location (i, j) within the CCD grid, Φ_p , is

$$\Phi_p = \int_x \int_y E_p(i, j) dy dx, \quad (5.9)$$

where x and y are the dimensions of the sensor site and E_p is irradiance incident on the site given by Equation 5.8. Assuming the surface patch in the scene, from which this flux originates, to be small and the surface to be within the camera's depth of field, it can be assumed that irradiance is constant over the area of the sensor, hence,

$$\Phi_p = A_p E_p(i, j), \quad (5.10)$$

where A_p is the area of the sensor.

The range of electro-magnetic radiation to which a CCD is sensitive is wider than that of the human visual system. Figure 5.2 shows the relative spectral response of a typical CCD device. It can be compared to the photopic equivalent for the human visual system. It can be seen that the CCD device is responsive to a broader spectrum and is particularly responsive to infrared radiation. This sensitivity can be reduced by the use of appropriate filters to block out unwanted radiation. The spectral response of a camera can be determined by the use of calibration methods (see Section 5.4) and a response function, $s(\lambda)$, can be formulated. A discrete pixel value V is

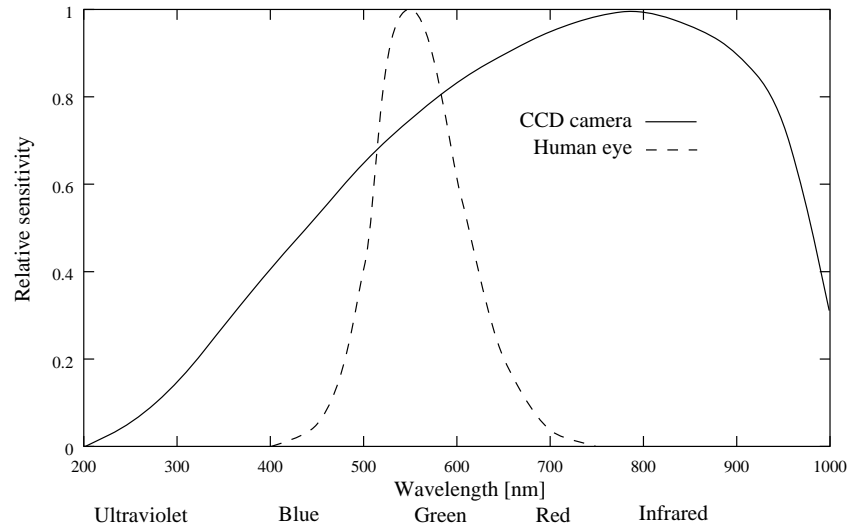


Figure 5.2: Relative sensitivity to wavelength of a CCD camera and the human eye.

obtained by quantisation of the total irradiance incident on the sensor site during the exposure time set by the shutter speed,

$$V = T Q A_p \int_{\lambda} s(\lambda) E_p(\lambda) d\lambda, \quad (5.11)$$

where T is the integration time as determined by the shutter and Q is a model of the quantisation process in the form of a *transfer function*. The transfer function defines the rate at which the pixel value increases in proportion to an increase in irradiance. This function also defines the dynamic range of the pixel values.

Commercial CCD cameras have primarily been designed for the acquisition of images for display purposes. This has significance when using these cameras for image processing in the context of measuring scene radiance. It should be noted that the transfer function of the camera may not be linear and hence an increase in scene radiance is not complemented by an proportional increase in image pixel value. This is due to the display *gamma* typically used to provide a qualitative representation of the scene. This can be compensated for by the use of gamma correction methods [36].

5.3.2 Colour Image Formation

In order that a colour image may be produced, the scene radiance is sampled in three wavelength regions representing red, green and blue light. This is typically performed using filters with trans-

mittance functions $\tau_R(\lambda)$, $\tau_G(\lambda)$ and $\tau_B(\lambda)$, see Figure 5.3. From the composition of these three samples the variation of wavelength from the scene may be approximated. The errors incurred by the use of such a tristimulus colour approximation are discussed in [36].




Figure 5.3: Relative sensitivity of CCD camera to filtered wavelengths.

The use of three filters to sample the scene radiance yields a triplet of values for each image pixel, (V_R, V_G, V_B) . The composite RGB pixel value can be obtained by extending Equation 5.11:

$$V_R = TQA_p \int_{\lambda} s(\lambda) \tau_R(\lambda) E_p(\lambda) d\lambda, \quad (5.12)$$

$$V_G = TQA_p \int_{\lambda} s(\lambda) \tau_G(\lambda) E_p(\lambda) d\lambda, \quad (5.13)$$

$$V_B = TQA_p \int_{\lambda} s(\lambda) \tau_B(\lambda) E_p(\lambda) d\lambda. \quad (5.14)$$

5.4 Sources of Error and Calibration

The processes by which the scene is projected and the image recorded are subject to a number of error sources. These errors originate from non-uniform projection by the optics of the scene and the sensor used to create the image. The detection and subsequent reduction of such errors can be achieved by the use of calibration methods. This section summarises the problems associated with CCD cameras and calibration techniques.

5.4.1 Optical Errors

Lenses do not provide a uniform projection across the area of the image. This results in two forms of aberration: spherical and chromatic. The use of spherical lenses, as opposed to the ideal parabolic shape, results in the scene being projected non-uniformly onto the image plane. The resulting image may exhibit blurring due to objects not being brought to focus uniformly across the image. Chromatic aberration is due to the refractive nature of the lens material. The extent to which light is refracted by the lens is a function of wavelength. Light, originating from the same scene point, of differing wavelength will not be focused to precisely the same point on the image plane. The aberration thus caused is particularly evident at the periphery of an image. Novak et al. suggest a method for the detection and correction of chromatic aberration [99].

Commercial lenses attempt to reduce aberrations. In practice camera lenses consist of a system of optics, typically a double-Gauss system. A physically-based model of such lenses is provided by Kolb et al. [69]. This model allows for the computation of irradiance incident on the image plane to be made for a variety of commercial lenses. Such a model is invaluable for the accurate radiometric analysis of image pixel values from cameras. Chromatic aberration may be reduced by the use of coatings applied to the surface of the lens.

Determining the perspective projection provided by the lens, and hence the parameters to the pin-hole approximation, can be achieved using a number of calibration methods. An overview of the various techniques is provided by Tsai [137].

5.4.2 Sensor Errors

The process by which a CCD chip generates an image of the spatial distribution of incident radiance is not without inherent errors. The silicon which makes up a CCD chip is thermally sensitive and may produce a signal due to ambient temperature. This results in *dark noise* across the image. The quantum process by which electrons are generated by incident photons is subject to uncertainty and results in *shot noise*. The dynamic range of a sensor site may cause regions of the image to be over or under exposed. Over exposure is a particular problem when using a CCD sensor since the resulting overflow of charge from a sensor site causes neighbouring sites to produce increased signals. The resultant *blooming* in the image causes highlight regions to appear

enlarged in the image. Detection and calibration of the errors due to the use of CCD sensors is presented by Healey and Kondepudy [44, 45] and Klinker [66].

The projection of the scene onto a segmented sensor to produce a digital image is the cause of spatial aliasing artifacts. The resulting errors, due to the quantisation of the projection process, are discussed in [40] and [57].

The sensitivity of the CCD sensors to both change in intensity and wavelength of light varies between manufacturers. For accurate analysis of the resulting pixels it is necessary to determine the transfer function of the sensor. Such calibration can be achieved by the use of test targets such as the Macbeth ColorChecker [84, 31]. Such targets provide samples of known chromaticity and reflectance so that the linearity of the camera's response can be assessed. Calibration methods for determining sensor sensitivity are provided by [66, 99, 146].

5.5 Summary

This chapter has briefly discussed the characteristics of CCD devices for the purposes of obtaining digital images from which radiometric measurements can be taken. The development of a camera model has shown that, given suitable calibration, such images can be used for measuring the reflected radiance from a surface patch.

Despite the inherent errors of such devices, CCD cameras and digitised photography may be used for surface reflection model estimation. The accuracy of the resulting model will be dependent upon the accuracy of the sensor used and of the sensor model. Errors in the measurements can be accommodated by a system but only if the magnitude and distribution of such errors has been determined by prior calibration.

Chapter 6

Surface Model Estimation

The aim of this chapter is to show how the reflection characteristics of an observed surface, illuminated by natural daylight, may be estimated. Chapter 3 has discussed how the reflection characteristics of a surface can be described in terms of a functional model. The most general description of reflection is given by the bidirectional reflectance-distribution function (BRDF) (see Section 3.1.2), which defines reflected radiance in terms of incident radiance confined within differential solid angles. Surface BRDF can be approximated for use in computer graphics and machine vision tasks using the reflection models presented in Chapter 3. It is the parameters of such a reflection model that are to be estimated by the methods presented here.

The appearance of an object is determined by both the illumination and the object's surface properties. This work assumes that the illumination consists of sun and sky light alone and corresponds to the models presented in Chapter 4. It is therefore necessary to determine whether, given *a priori* knowledge of the illumination and a sensor measurement of reflected energy due to this illumination, the parameters of a reflection model can be determined. Further to this, it is necessary to determine whether a solution is unique and to what degree of certainty the solution found is accurate.

The use of the sensor model developed in Chapter 5 allows image pixel values obtained from a camera to be interpreted in the context of a surface reflecting incident radiance in the direction of a sensor. In a natural scene, this illumination is due to sun and sky light. These are characterised by the models of Chapter 4. This work only considers regions of the scene which are within the

camera's *principle plane of focus*. Only these objects are brought to sharp focus on the image and hence, for the purposes of this work, it can be assumed that radiance from a small area on an object's surface is projected onto the pixel area of the sensor plane [51]. Surface interreflection is not considered.

An image of a scene represents a number of measured values. Each value is the result of reflected light in the direction of the camera. Given accurate models of daylight illumination, surface reflection and sensor response it is possible to predict these measured values for any given scene. Any discrepancy between the measured and predicted values would be due to either measurement error or inadequacies in the models. Given *a priori* knowledge of the scene geometry, the sensor, and the nature of the illumination, the free parameters of the complete model for predicting pixel values are those which describe the surface reflection characteristics. We wish to select these parameters such that the discrepancy between measured values and those predicted by the model are minimised. This chapter presents a measure of difference between such values and methods by which it can be minimised.

It should be noted that a surface's BRDF cannot be measured directly. This is due to the differential quantities used in the definition of BRDF. As such, any model of surface BRDF derived from finite measured data is only an approximation. As Nicodemus states [95];

‘The BRDF itself, as a ratio of infinitesimals, is a derivative with “instantaneous” values that can never be measured directly. Real measurements involve non-zero intervals of the parameters, e.g $\Delta\omega$ or $\Delta\lambda$ rather than $d\omega$ or $d\lambda$, and hence, can yield only average values f_r over these parameter intervals.’

6.1 The χ^2 Metric

The difference between a set of measured values and those predicted by a model can be evaluated in a *least-squares* sense. This provides a figure-of-merit function which evaluates the correspondence between measured data and a model. Given a set of m observations, y_i , and an associated set of model values, $y(x_i; \mathbf{a})$, obtained for the same data points, x_i , the least squares difference between

the data and the model is given by

$$\text{least-squares difference} = \sum_{i=1}^m [y_i - y(x_i; \mathbf{a})]^2, \quad (6.1)$$

where x_i are the model variables and \mathbf{a} is a vector of parameters upon which the model depends. Selection of these parameters, such that the least-squares difference is minimised, provides the model which most accurately describes the given data.

In practice measured data will have an associated error. In the case of digital images, the sources of such error are described in Chapter 5. With suitable calibration the error of a sensor may be determined and knowledge of possible measurement error may be included in the figure-of-merit function. If each measured data value, y_i , has an associated Gaussian error with standard deviation σ_i , a *chi-squared* metric may be used [5]:

$$\chi^2 = \sum_{i=1}^m \left[\frac{y_i - y(x_i; \mathbf{a})}{\sigma_i} \right]^2. \quad (6.2)$$

Minimisation of this function will yield the parameters which best model the measured values given the estimate of the measurement errors. The process by which these optimal parameters are determined is termed *regression*.

Minimisation of χ^2 to find the optimal model parameters is a method of *maximum likelihood*. Given a model, it is assumed that the set of measured data values are observations from the parent distribution of possible model values. Minimising χ^2 is the equivalent of maximising the probability that the observations are from the parent distribution. The performance of the least-squares method of model parameter estimation is governed by the accuracy of the sensor readings and the quality of the model. It is also dependent upon the number of observations made. Measured values which are not within the error estimation σ_i cause problems when considering the χ^2 metric. Such outlying values in the measured data will give a poor fit between data and model and yield poor estimates for the parameters.

It can be seen from Equation 6.2 that χ^2 is a function of the model parameters. The χ^2 function therefore describes a hypersurface in a space of dimension equal to the number of model parameters. The global minimum of this hypersurface is the point at which the parameter values provide a least-squares fit with the measured data. Minimisation of Equation 6.2 is equivalent to locating the global surface minimum within the parameter space.

6.2 Description of Data and Associated Model

Consider the simple case of a surface illuminated by a single point source with reflected energy being measured at a single discrete wavelength by a sensor. The intensity of the reflected energy at this wavelength is due to the intensity of the point source, the reflectance of the surface, and the extent to which light is scattered and attenuated by the surface structure. Surfaces which exhibit different levels of scattering and absorption may appear similar to a sensor which is only receiving reflected light from a single, small direction, see Figure 6.1.




Figure 6.1: *Sensor measuring reflected radiance within a small solid angle due to a point source: (a) Lambertian surface; (b) glossy surface. In each case the magnitude of the reflected radiance incident on the sensor is the same.*

This simple example highlights the problems of the task. The sensor can provide no information regarding light reflected into directions other than those which are incident on the sensor. A single measurement provides information relating only to that particular instance of the viewing and illumination geometry. Using the single measurement from a calibrated device would not provide suitable information with which to determine the surface's reflection characteristics.

So that more information can be obtained regarding the surface's reflection characteristics it is necessary to obtain further measurements. In this monochromatic case, the BRDF of a surface is a function of the incident and reflection directions relative to the surface normal. Hence, it would be advantageous to obtain a number of measurements each of which is a result of a different illumination or viewing geometry. Each measurement provides data to which a proposed reflection model can be compared. Figure 6.2 shows possible methods of obtaining further reflection measurements.

Moving the light source, as shown in Figure 6.2(a), will provide a number of measure-




Figure 6.2: *Methods for obtaining further measurements of surface reflection: (a) moving the light source; (b) moving the sensor; (c) sampling other regions of the surface.*

ments from which the surface reflection model can be estimated. In laboratory conditions the light source can be positioned such that many measurements can be obtained from a wide distribution of source positions. In the case of a scene illuminated by natural light alone, such control over light source position is not possible. However, if the scene is observed over a period of time, the relative position, and hence surface irradiance, of the sun will change, as will the spatial and spectral distribution of radiance from the sky. These temporal changes are characterised by the models given in Chapter 4 and hence provide a variety of illumination conditions from which measurements can be obtained. In the case of a naturally illuminated scene, light is incident from a possible hemisphere of directions, depending upon the surface orientation. Each of the directions may contribute to the light reflected in the direction of the sensor.

Figure 6.2(b) shows that further measurements can be obtained from a single surface patch by moving the position of the sensor, hence changing the viewing geometry. This method is not considered here.

The use of an image obtained by a CCD camera provides a number of measurements, each of which may be considered to be due to light reflected from a different surface patch. As shown in Figure 6.2(c), each pixel of such an image represents a different illumination and viewing geometry. The extent to which these geometries vary is dependent upon the change of surface normal across the visible surface.

It is proposed that a combination of changing illumination and the sampling of different surface regions will provide data suitable for a surface's reflection model to be determined. The basic approach will be to observe a surface over a period of time, thus providing varying illumination

conditions. The temporal change in value of a single pixel across the resulting image sequence provides a measure of surface reflection from that surface patch, given the changing illumination. Since the projection of the visible surface may cover a group of image pixels, a number of such measurements can be obtained from each individual image in the sequence. Each pixel in an image sequence represents a measured data value, y_i , to be used in the evaluation of Equation 6.2. Figure 6.3(a) shows the change in value of a single pixel across an image sequence observing a glossy surface under a clear sky. The variation in the values across the sequence is due only to the change of illumination over the duration of the sequence. The extent of this variation is determined by the reflection characteristics of the surface. Figure 6.3(b) shows the values of a number of pixels from the same surface image. Here it can be seen that, though the material has remained constant, the differing illumination and viewing geometries presented by each pixel has resulted in a variation in values between the pixels. Again, the extent to which the values from different pixels vary will be determined by the reflection characteristics of the surface.




Figure 6.3: *The change in image pixel value across a sequence of thirteen images for: (a) a single pixel in each image; (b) three pixels in each image.*

6.2.1 Selection of a Suitable Model

The model, $y(x_i; \mathbf{a})$, used in Equation 6.2 is to be defined such that it describes the variation in observed image pixel values over an image sequence. The variables, x_i , of this model represent the camera position, time the image was taken, and the orientation of the surface. From these variables the illumination due to sun and sky light can be determined, as can the direction of reflection towards the camera. The unknown parameters of the model, \mathbf{a} , are therefore those which describe the surface's reflection characteristics such as roughness, Fresnel effects or wavelength

attenuation.

Initial work proposed the use of a surface reflection model alone to describe the variation of image pixel values [78]. In this case, illumination from the sky was neglected and the sun was modelled as a moving point source. The surface reflection model used was a combination of those proposed by Lambert and Phong (see Chapter 3). Results showed that, whilst this model of pixel value variation across a sequence did allow the estimation of model parameters, there were strict limitations as to the success of the method. Since only direct solar illumination was considered it was necessary for the surface to be orientated such that it faced the sun. This limitation obviously excluded the consideration of overcast skies. The results were also shown to be dependent upon the location of the camera with respect to the surface. In this case the model chosen for $y(x_i; \mathbf{a})$ did not attempt to predict image pixel values. It attempted only to describe the variation of values across the sequence and not absolute pixel values.

The use of the full sensor model presented in Chapter 5 ensures a more accurate description of pixel value variation. Equations 5.8 and 5.12 show how daylight illumination and surface reflection models can be incorporated into the sensor model so that pixel values may be predicted. The use of this more complex model allows illumination due to sky light and the characteristics of the sensor to be accounted for. It is intended that the provision of this, more accurate, model will provide improved results over those presented in [78].

It is proposed to evaluate χ^2 using image pixel values and pixel values predicted by the sensor model of Chapter 5. Given *a priori* knowledge of the illumination and scene geometry, the free parameters of this model are those of the function used to describe surface reflection. All other aspects of the sensor model shall be assumed to have been recorded at the time the image was taken. Therefore, the evaluation of $y(x_i; \mathbf{a})$ will be performed using the full sensor model with \mathbf{a} being the parameters of the chosen surface reflectance function. In this way, the optimal parameters of the reflectance function may be estimated by minimisation of the χ^2 function for a given image sequence.

6.2.2 χ^2 for an Image Sequence

Given that the supplied data are to be in the form of a sequence of images, the χ^2 function can be redefined specifically for this application. Considering a single pixel across a sequence of m

images, I , the value of the pixel, in a single image, with co-ordinates x, y is denoted by $I^i(x, y)$ where $i = 1, \dots, m$. A single pixel of a colour image is typically represented by a triplet of values. These represent the sensor response to three wavelength regions (see Section 5.3.2). Each of these values may be treated in isolation and is a result of sampling the model over the range of wavelengths to which light has been filtered. The tristimulus value of a single pixel in a colour image of the sequence is, therefore, $\{I_R^i(x, y), I_G^i(x, y), I_B^i(x, y)\}$, where R, G, B denote the three sampled wavelengths represented by each of the values.

The model used to describe the variation of pixel values across an image sequence is discussed in Section 6.2.1. Again, in the case of colour images, three values may be associated with a pixel. The tristimulus pixel value predicted by the model for a given set of reflection model parameters \mathbf{a} is denoted by $\{M(i, x, y, R; \mathbf{a}), M(i, x, y, G; \mathbf{a}), M(i, x, y, B; \mathbf{a})\}$, where M is the pixel value model. Note that the image number, i , the pixel co-ordinates, and the filtered wavelengths are variables of the model. These denote the time the image was taken, the sampled wavelengths of light, and the location of the pixel within the considered image.

If the values of a number of colour pixels in each image of a sequence are considered then;

$$\chi^2 = \sum_{i=1}^m \sum_{x,y} \sum_{\lambda} \frac{1}{\sigma_i^2} [I_{\lambda}^i(x, y) - M(i, x, y, \lambda; \mathbf{a})]^2, \quad (6.3)$$

where $\lambda = \{R, G, B\}$, x and y are the co-ordinates of each considered pixel, and σ_i is the expected error associated with the pixel value.

In the above definition of the χ^2 function, the difference between each individual pixel and that predicted by the model is summed individually and independently of other pixel values in the sequence. In the case of colour images, each of the RGB values is treated as an individual measurement. Therefore, a single evaluation of the term to be summed only calculates the ability of the model to predict that particular pixel value. Only by the use of the complete summation can the ability of the model to predict values across the whole sequence be determined.

Equation 6.3 is a good measure of the difference between image and predicted pixel values across an image sequence since, in its evaluation, each frame, and each pixel within the frame, is considered independently. If a set of measured pixel values and predicted values have the same combined magnitude but the values occur in a different order in the sequence then the χ^2 value will be high. To achieve a low χ^2 value both the magnitude and order of the pixel values must be similar. Therefore, both the shape and the magnitude of the curves presented in Figure 6.3 must

be modelled by the function selected for $M(i, x, y, \lambda; \mathbf{a})$ in Equation 6.3.

When dealing with digital images, pixel values are typically integer quantities. Therefore the calculation of difference between actual and predicted pixels in Equation 6.3 involves integer values (though predicted values may be estimated to floating point accuracy). It is also the case that, given an accurate model of the imaging process, as the estimated parameters approach those of the observed surface, the difference in Equation 6.3 will tend to zero. Both of these situations give a potential for error in the calculation of χ^2 for an image sequence. However, as shall be shown in the results (see Chapter 7), in the case of real image sequences, χ^2 rarely approaches zero and in some cases may be considered large.

As stated in Chapter 5 the error in the actual pixel values is due to sensor and optical error. In the case of CCD devices this error is a quantum process and is therefore more adequately modelled as Poisson distribution. The magnitude of optical error varies with respect to image pixel coordinates and is not uniform across the image. It is therefore an approximation to model these expected errors with the single Gaussian term σ_i^2 in Equation 6.3. For these reasons the use of the χ^2 metric in the case of digital images may not be seen as ideal. The derivation of a more suitable metric is seen as a possible direction for future work.

6.3 Topology of the χ^2 Hypersurface

As stated in Section 6.1, the optimal model parameters are to be estimated by minimisation of the χ^2 function. To determine the most appropriate method for searching the parameter space, the expected topology of the χ^2 hypersurface is now considered. Only with an understanding of the nature of the χ^2 function can methods be developed which will allow minimisation.

Synthetic image sequences provide noiseless data obtained using a precise camera model. To generate these image sequences the *Radiance*¹ lighting simulation system has been used. The images are of a scene composed of an object with known surface model parameters and use the sky models presented in Chapter 4. We aim to search the parameter space of the same reflection model used to generate the image sequence, and thus would expect to find a χ^2 value of zero where the reflection model parameters are the same as those used to generate the image sequence.

¹See Section 4.2.4 for a description of this system.

The reflection model used in the *Radiance* lighting simulation system is that due to Ward [141] and has been described in Section 3.4.5. Sampling light at three discrete wavelengths, representing red, green and blue light, provides an approximation of the spectrum of reflected light. In this case, the colour reflection model has five parameters which represent the surface reflectance at three discrete wavelengths, surface roughness and a measure of surface gloss. Visualisation of the resulting five-dimensional χ^2 hypersurface in the parameter space, for a given image sequence, is not possible. Reducing the number of free reflection model parameters to two will allow the surface to be presented. To allow this, the two-parameter version of this reflection model described in Section 3.4.5 shall be used. Here, the two free model parameters represent surface gloss, ρ_s , and surface roughness σ_α . All other parameters are fixed at appropriate values. Observations made from the examples in this reduced parameter space may not necessarily extend to the full reflection model, or to other reflection models. However, minimisation within a higher dimensional parameter space is expected to be at least as difficult as these observations will show.

To appreciate the topology of the χ^2 surface for a specific image sequence we first consider a horizontal, unoccluded Lambertian surface ($\rho_s = 0$, $\sigma_\alpha = 0$), illuminated by a sun and clear sky, and observed from a camera situated in the north such that it looks down upon the surface at an angle of 45° to the surface normal. The image sequence consists of thirteen images, each taken hourly between 06:00 and 18:00. Selected images from this test sequence are shown in Figure 6.4.




Figure 6.4: Selected images from a sequence showing a Lambertian disc, illuminated by a clear sky, on June 21 at terrestrial location $40^\circ N, 0^\circ W$. The camera is situated in the north and looks down upon the disc at an angle of 45° . Each image is taken at (a) 08:00, (b) 12:00, (c) 14:00, (d) 18:00.

First the χ^2 function for a single pixel at the centre of each image is considered. Note that the χ^2 value at each point in the parameter space is found by evaluating Equation 6.3 for the same pixel within each image of the sequence. Figure 6.5 shows the χ^2 surface for this pixel within the permitted parameter space of the Ward model. Note that, as expected, the minimum in the

parameter space occurs on the line $\rho_s = 0^2$.




Figure 6.5: χ^2 function for a single pixel in an image sequence of a horizontal Lambertian disc illuminated by a clear sky and viewed from the north.

Figure 6.6 shows the χ^2 function for an image sequence, of the same length and duration, observing a disc generated with various reflection model parameter values. In each case the minimum χ^2 coincides with the model parameters used to generate the original sequence. These examples show that, in the situations considered, a single minimum exists in the χ^2 function. This minimum coincides with that of the optimal parameter values in each case. All other parameter value combinations yield a higher χ^2 value. It is therefore proposed that the minimisation of the χ^2 function will lead to correct estimation of the model parameters in these cases.

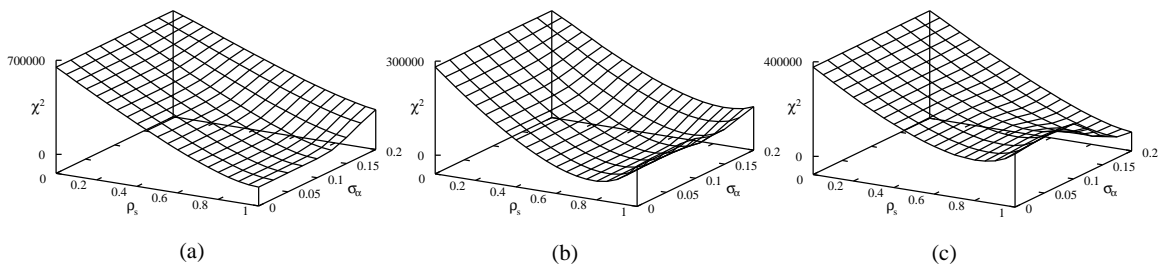


Figure 6.6: χ^2 function for a single pixel in an image sequence of a horizontal disc illuminated by a clear sky and viewed from the north: (a) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (b) disc with Ward parameters $\rho_s = 0.8, \sigma_\alpha = 0.1$; (c) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

6.3.1 Effect of Camera Position

We now consider the same scene and illumination conditions as above but viewed from differing camera positions. It would be expected that the existence of a highlight due to a bright sun would

²Surface roughness, σ_α , has no significance in the Ward model when $\rho_s = 0$.




Figure 6.7: χ^2 function for a single pixel in an image sequence of a horizontal disc illuminated by a clear sky and viewed from the east: (a) disc with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

provide evidence that the observed surface is glossy. In the case of a camera sited in the south, the illumination geometry would be such that a highlight could not occur when using Ward's reflection model. In the following experiments, a camera in the east and south is considered. Since motion of the sun, and hence the spatial distribution of sky radiance, is symmetric about the north-south axis, it is not necessary to consider a camera in the west. The effect of other camera positions can be inferred from these examples.

Figure 6.7 shows the χ^2 function values generated by the east view sequence for a variety of surface types. This surface shows a more complex topology than that shown in Figure 6.6. The ridged nature of the surface shows that a gradient descent search method (see Section 6.4) would be inefficient in this case. In general, the global minimum is surrounded by steep regions. Away from the optimal parameter values, plateau and local minima exist in these surfaces. Such regions on the χ^2 surface may cause difficulties for a method which searches the parameter space for the optimal parameters.

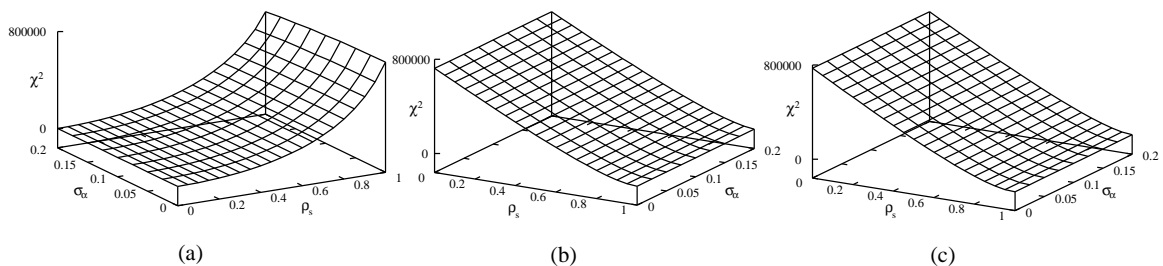


Figure 6.8: χ^2 function for a single pixel in an image sequence of a horizontal disc illuminated by a clear sky and viewed from the south: (a) disc with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

Figure 6.8 shows the χ^2 function for scenes observed from the south. As stated above, with the Ward reflection model a glossy highlight would not be expected to be within the sequence with the camera in this position relative to the surface. It can be seen from the graphs that the resulting topology lacks features, though surface gloss has been correctly indicated by a low χ^2 function value. However, changes in surface roughness are not observable due to the absence of any highlight information. The χ^2 surfaces of Figures 6.8(a) and 6.8(b) do not show any variation with surface roughness.

The above experiments show that the position of the camera has significant effect on the expected performance of the system. The absence of highlight information, such as that missing from the data obtained from the south view, prevents determination of surface roughness. In the absence of this information any system based on χ^2 minimisation would be unable to determine whether it is observing a Lambertian surface of low albedo or a glossy surface with a geometry such that a glossy highlight is not visible in the sequence. Both the north and east camera positions provide some highlight information due to the movement of the sun over the duration of the image sequence. The above examples show the significance of the sun for the estimation of surface parameters.

6.3.2 Effect of Sky Conditions

A CIE clear sky, as considered in Section 6.3 and Section 6.3.1, provides a hemispherical illuminant which has temporally varying spatial distribution of radiance over the duration of an image sequence. It also allows direct solar illumination, due to the sun, to be considered. It is therefore expected that such illumination would provide the most suitable conditions for reflection model parameter estimation. This section considers scenes which have been illuminated by alternative sky conditions such as those described in Chapter 4.

Figure 6.9 shows the χ^2 function for scenes illuminated by an intermediate sky ($\phi = 0.5$) and observed by a camera in the east position. These surfaces are similar to those of Figure 6.7. However, the overall gradient of the surface is reduced, which indicates a lessening of the difference between the possible pixel values predicted by the model. Figure 6.9(b) clearly shows a local minimum which has a significantly higher χ^2 value than that of the global minimum. This may present problems to a search algorithm which is based upon a gradient descent of the surface.

It is expected that, since these examples show a similar topology to those presented for a clear sky, the results obtained from intermediate skies will be comparable.




Figure 6.9: χ^2 function for a single pixel in an image sequence of a horizontal disc illuminated by an intermediate sky and viewed from the east: (a) disc with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

Figure 6.10 shows the χ^2 function for scenes illuminated by a CIE overcast sky and observed by a camera in the east position. In comparison with results presented for the same scenes illuminated by clear and intermediate skies, these surfaces have low χ^2 values. This indicates that there is a much reduced difference between the pixel values predicted for differing model parameters. However, the topology of the surface in each case is not featureless. It can be seen that scenes of differing surface types have produced strikingly similar χ^2 surfaces. Only in the cases shown in Figures 6.10(a) and 6.10(b) does the surface slope toward the optimal parameters. However, in each of these cases the minimum covers a region of the surface and as such a unique optimal solution could not be determined. In the case shown in Figure 6.10(c) minimisation of χ^2 would lead to incorrect parameter estimation.

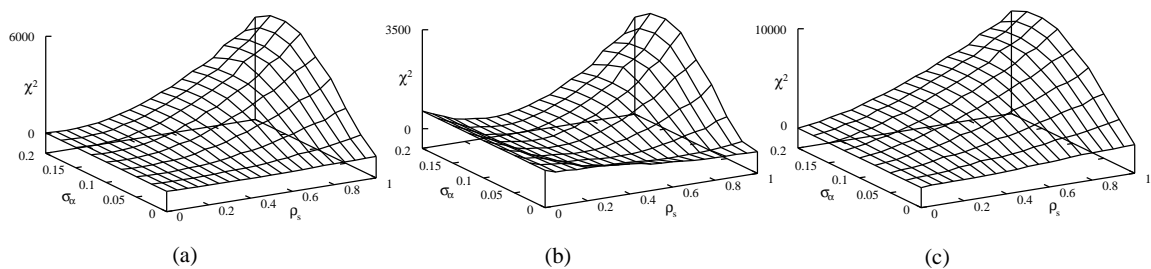


Figure 6.10: χ^2 function for a single pixel in an image sequence of a horizontal disc illuminated by an overcast sky and viewed from the east: (a) disc with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

6.3.3 Effect of Camera Noise

The expected effect of camera noise on the system's performance is now considered. As described in Chapter 5, any practical image acquisition system is subject to noise. In the case of CCD digital cameras, noise comes from a variety of sources such as optics, sensor and quantisation. For the purposes of system evaluation camera noise is modelled here as an additive Gaussian process. Whilst this is not an accurate representation of the expected noise, it serves to provide a measure of system performance. Care must be taken in using additive noise so as to not exceed the actual dynamic range of the pixel values. A negative pixel value, for example, is not permissible. In these examples, the standard deviation of the Gaussian noise is expressed as a percentage of the pixel dynamic range.

Figure 6.11 shows the value of a pixel throughout an image sequence viewing a rough metallic disc, illuminated by a clear sky and observed by a camera positioned in the east. The graph shows the values of the same pixel for clean images which contain no noise, images to which Gaussian noise with a standard deviation of 5% has been added, and images to which Gaussian noise of with a standard deviation of 10% has been added.




Figure 6.11: Value of a pixel in each image of a sequence to which Gaussian noise has been added.

Figure 6.12 shows the χ^2 function for the pixel over the sequence to which Gaussian noise with a standard deviation of 5% has been added. These should be compared with those in Figure 6.7. It can be seen that, in this case, the general topology of the surfaces has been unaffected by the addition of noise. The effect of noise has been to reduce the values of the χ^2 function for any set of parameter values, and hence reduce the gradient of the function.




Figure 6.12: χ^2 function for a single pixel in an image sequence of a horizontal disc illuminated by a clear sky and viewed from the east. Additive Gaussian noise with a standard deviation of 5% has been added to each image: (a) disc with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

Figure 6.13 shows the χ^2 function for the same sequence but with Gaussian noise with a standard deviation of 10%. Again, the topology is similar to that found in the clean images shown in Figure 6.7. These examples serve to show that the χ^2 function is suitable for the comparison of pixel values and model predictions in the case of noisy images. The general topology of the examples has remained similar to those presented for clean images. The effect of noise has been to reduce the χ^2 function values and hence the gradient of the function. As noise is increased it is expected that the χ^2 surface will flatten such that minimisation will not be possible. It should be noted that, as stated in Section 6.1, noise in the form of outliers compromises the least-squares fit. Such outliers may, in the context of images, be incorrectly interpreted as highlight peaks in the sequence values.

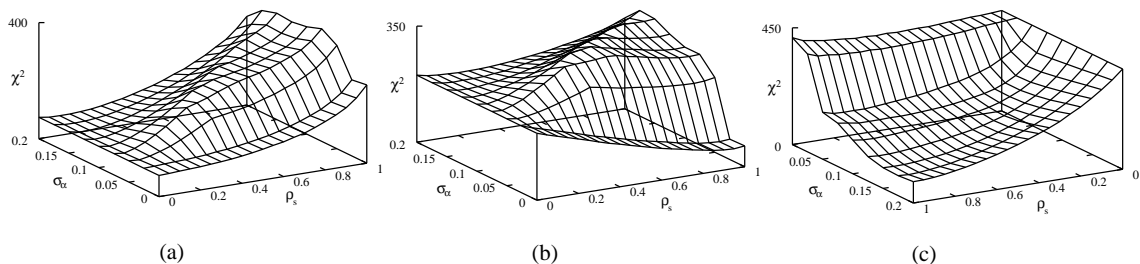


Figure 6.13: χ^2 function for a single pixel in an image sequence of a horizontal disc illuminated by a clear sky and viewed from the east. Additive Gaussian noise with a standard deviation of 10% has been added to each image: (a) disc with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

6.3.4 Sampling Multiple Pixels with Similar Surface Normals

It would be expected that sampling more than one pixel of the 2D projected surface image would provide further evidence to support a set of hypothesised reflection model parameters. This section considers the same image sequences as those presented above but with four pixel values used in the calculation of the χ^2 function.

Figure 6.14 shows the value of four pixels across the image sequence for a number of possible reflection parameter combinations. Figure 6.14(a) shows the pixel values sampled from a Lambertian disc. As expected the slight variation in viewing geometry provides no difference between the pixel values. Values sampled from a smooth metallic disc, shown in Figure 6.14(b), show slight variation. The peak value for each pixel varies, though the small highlight has not been sufficiently sampled by these pixels. Figure 6.14(c) shows pixel values sampled from a rough metallic disc. It can be seen that the peak pixel value occurs in a different image of the sequence in each case. In this case however, the increased size of the highlight, due to the rough nature of the surface, has meant that it appears in the graphs for these pixels.

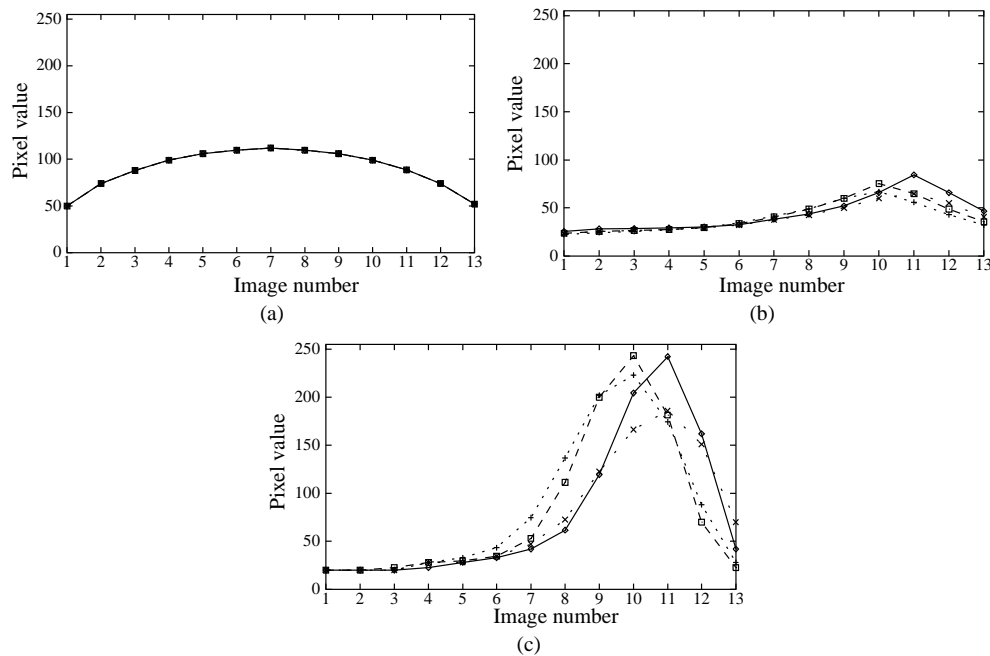


Figure 6.14: Values of four pixels within an image across a sequence: (a) disc with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

Figure 6.15 shows the χ^2 function for these scenes. These should be compared to those of Figure 6.7 which show the same sequence but with a single pixel being used for the calculation of χ^2 . It can be seen that the gradient of the function is higher than those for a single sampled pixel. This is to be expected due to the increased number of summations performed in evaluating Equation 6.2. The addition of further pixels also has the effect of smoothing the χ^2 function in each case. This would be advantageous for any search method that is to be used to find the optimal parameters.




Figure 6.15: χ^2 function for four pixels in an image sequence of a horizontal disc illuminated by a clear sky and viewed from the east: (a) disc with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

6.3.5 Sampling Multiple Pixels with Differing Surface Normals

Figure 6.16 shows selected images from a sequence observing a Lambertian sphere illuminated by a clear sky and viewed from an east position. The sequence consists of thirteen frames taken hourly, starting at 06:00 and ending at 18:00. The change in surface normal provides a variation of shading across the surface of the sphere. It is expected that sampling a number of pixels across the image sequence would provide suitable variation of pixel values so that the reflection model parameters may be determined.

Pixel values sampled from sequences imaging spheres of differing surface models are shown in Figure 6.17. In each case the values of five pixels are shown. Figure 6.17(a) shows values sampled from a sequence of a Lambertian sphere. Owing to the change in surface normal there is variation in the pixel values. This is to be compared to the values sampled from a Lambertian disc shown in Figure 6.14, where no variation can be observed. Figure 6.17(b) shows pixel values sampled from a smooth metallic sphere. The selected pixels have not adequately sampled the highlight, which is present due to the sun, hence the pixel values have little variation across the




Figure 6.16: Selected images from a sequence showing a Lambertian sphere, illuminated by a clear sky, on June 21 at terrestrial location $40^\circ\text{N}, 0^\circ\text{W}$. The camera is situated in the east and looks down upon the horizontal plane at an angle of 45° . Images taken at (a) 08:00, (b) 12:00, (c) 14:00, (d) 18:00.

sequence. Figure 6.17(c) shows values obtained from a rough metallic sphere. It can be seen that, in this case, the roughness of the surface is such that a highlight due to the sun has been sampled by some of the pixels. There is significant difference between the location of these peaks within the image sequence for each pixel.

The χ^2 surface for three image sequences of a sphere illuminated by a clear sky and observed from the east is shown in Figure 6.18. These show a similarity to those of Figure 6.7 and Figure 6.15.

The effect of sampling pixels of an object which represent regions of different surface normals for the problem situations highlighted above is now considered. Section 6.3.1 and Section 6.3.2 have shown that a camera situated in the south and objects illuminated by overcast skies presents problems for the estimation of surface model parameters using χ^2 minimisation. Figure 6.19 shows the χ^2 surfaces for spheres illuminated by a clear sky and observed in the south. In comparison with the results shown in Figure 6.8, it can be seen that there is more detail contained in the surfaces. In the case of the metallic spheres ($\rho_s = 1$) the difference in surface roughness can now be observed. The addition of data from regions with different surface normals has, in this case, provided suitable information for greater accuracy to be achieved.

Section 6.3.2 highlighted the difficulties presented by an overcast sky. In this situation there is no illumination provided by the sun and sky light radiance changes only in magnitude and not spatial distribution. Figure 6.20 shows that the χ^2 surface has been smoothed with the sampling of further pixels. However, in the case of the rough metallic sphere, the minimum is not in the expected position of $\rho_s = 1$ and $\sigma_\alpha = 0.2$. This would result in the incorrect reflection model

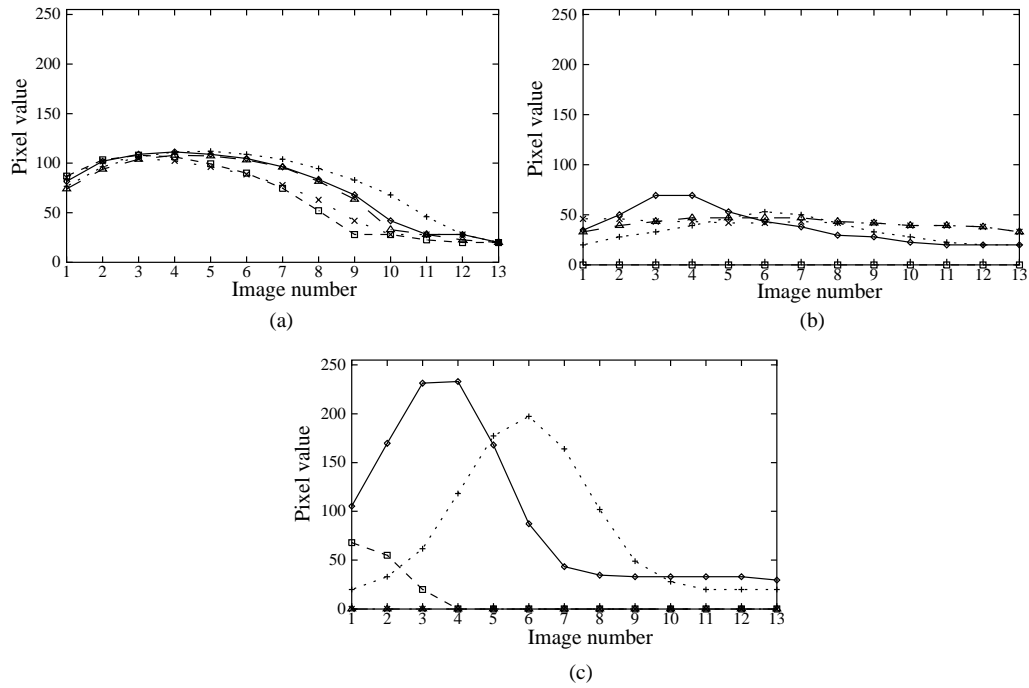


Figure 6.17: Values of five pixels within an image across a sequence: (a) sphere with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) sphere with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) sphere with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.




Figure 6.18: χ^2 function for five pixels in an image sequence of a sphere illuminated by a clear sky and viewed from the east: (a) sphere with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) sphere with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) sphere with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

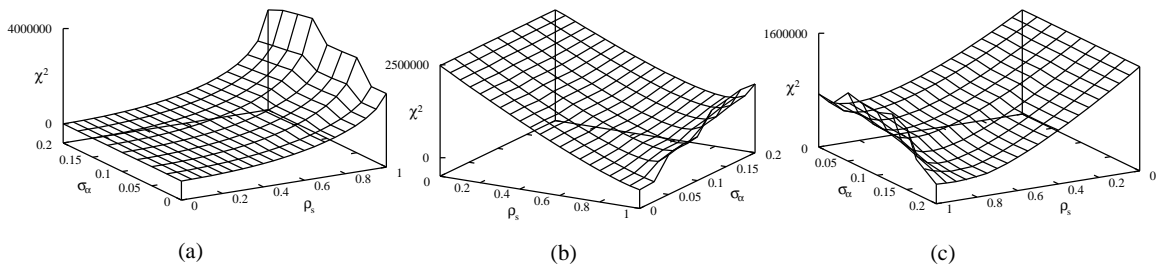


Figure 6.19: χ^2 function for five pixels in an image sequence of a sphere illuminated by a clear sky and viewed from the south: (a) sphere with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) sphere with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) sphere with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

parameters being estimated by minimisation of χ^2 in this case. This example serves to show the difficulties presented by overcast sky illumination for reflection model estimation.




Figure 6.20: χ^2 function for five pixels in an image sequence of a sphere illuminated by an overcast sky and viewed from the east: (a) sphere with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) sphere with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) sphere with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

6.3.6 Effect of Sequence Length and Image Frequency

The above sections have considered image sequences which cover the majority of the daylight period of a day and sampling the scene hourly. This section considers the effects of sequence length and image frequency on the possible performance of the system. Here image sequences of a disc illuminated by a clear sky and observed from the north are considered. Figure 6.21 shows the χ^2 surfaces for discs of a variety of materials observed over a sequence of five images taken hourly between 12:00 and 16:00. These should be compared with those of the complete image sequence shown in Figure 6.5 and Figure 6.6.

It can be seen that, owing to the short sequence length, and hence reduced number of sum-

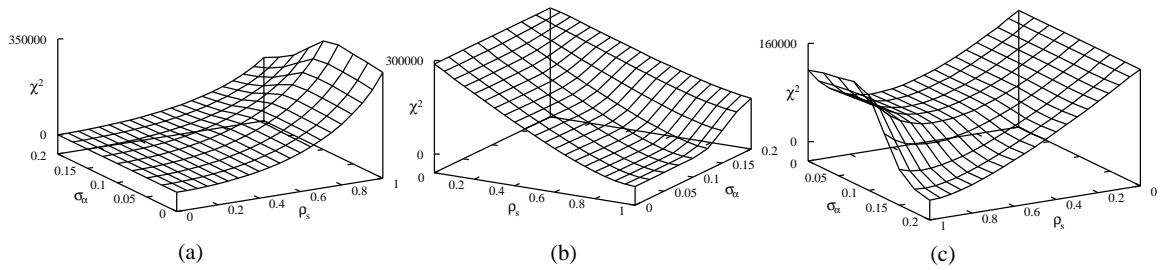


Figure 6.21: χ^2 function for one pixel in a short image sequence of a disc illuminated by a clear sky and viewed from the north: (a) disc with Ward parameters $\rho_s = 0.0, \sigma_\alpha = 0.0$; (b) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.0$; (c) disc with Ward parameters $\rho_s = 1.0, \sigma_\alpha = 0.2$.

mations performed in the evaluation of the χ^2 function, the gradient of the function is reduced. The overall topology of the surfaces is similar to those of the full sequence. Performance of the system would therefore be expected to degenerate as the sequence length becomes shorter. Fewer pixel samples across an image provide less supporting evidence for a particular set of reflection model parameters. It would also be expected that, if the period over which the sequence is taken is such that the illumination and reflection geometry does not allow the characteristic highlights of glossy surface to be observed, then the performance of the system will be poor. For example, a camera sited in the east would not observe a highlight due to the sun in a sequence of images taken in the morning hours of a day. This is a similar problem to that encountered with the camera sited in the south. In such a situation the system will have difficulty in differentiating between Lambertian and glossy surfaces using χ^2 minimisation.

The frequency with which images are taken of the scene will also affect the ability of a system to estimate surface reflectance. It has been found that images taken hourly or half hourly provide a suitable sampling frequency so that glossy highlights may be observed. If this frequency is reduced then the rate of change in illumination may be such that highlights, characteristic of some surface types, may not be observed in the pixel values of the image sequence. This dependence on frequency can be seen in Figure 6.22. Here the temporal change in surface reflection has not been sufficiently sampled over the duration of the sequence and the peak reflected radiance has been missed. A number of reflection models, characterising very different materials, may provide a good fit with the resulting data.




Figure 6.22: Sampling reflected radiance at three hour intervals, indicated by dashed lines.

6.3.7 Observations from Simple Experiments

The experiments presented above have shown the difficulties of determining surface reflection properties from a sequence of images. In each case the problem presented has been simplified so that the topology of the χ^2 function can be visualised. They have, however, served to highlight the conditions of illumination or viewing geometry which may cause problems for any system based upon the minimisation of the χ^2 difference between image pixel values and those predicted by a model.

The extension to a full colour reflection model would increase the dimension of the parameter space from that shown in the above examples. These extra parameters would account for surfaces of differing reflectance, Fresnel effects and surface colour. It is expected that the inclusion of these extra parameters would cause the topology of the surface to be at least as complicated as the above examples have illustrated. In these situations it is important that the method used to find the optimal parameters (those that yield a low χ^2 value) is efficient and robust. Given the continuous range of each of the parameters there is a potentially high number of parameter possibilities, hence the search space can be considered large.

It should be noted that the performance of the system is intrinsically dependent upon the performance of the camera. A camera which is not sufficiently sensitive to respond to the subtle changes, which a reflection model parameter attempts to model, will not provide suitable information for the system. Performance is also determined by the accuracy of the illumination and surface reflection models used. Natural surfaces and actual skies present a wide variety of observable effects. Some of these are not encapsulated in the models presented in Chapter 3 and Chapter

4.

6.4 Minimising χ^2

Section 6.3 has shown the topology of the χ^2 surface in the strictly limited case of a reflection model with two parameters. Observations made from the χ^2 examples show that the surface may contain local minima and plateau regions. In the case of a complete reflection model, the parameter space is multidimensional and hence the regression method chosen for this purpose must be suitable for such a search space. It is therefore necessary that the method used to minimise the χ^2 function should not be misled by plateau regions or non-global minima. Typically, a model of surface reflection has a non-linear dependence on the parameters. It is therefore the case that the complete sensor model, which utilises the surface model, is non-linear. Any regression method used to minimise χ^2 in this case must therefore be suitable for non-linear dependence of model parameters.

6.4.1 Brute Force Search

The simplest method of locating the parameters which yield a minimum in the χ^2 hypersurface is that of a brute force search. The parameter space can be sampled at regular increments of each of the parameters, Δa_j , comparing the function value with that of the minimum χ^2 thus found. Whilst this approach is certain to sample the parameter space with uniform density, it is not certain that it will locate the optimal solution. The global minimum may lie in between the sampled positions in the parameter space. The accuracy of this systematic search method is determined by the size of the increments used to sample the search space. Reducing the size of these increments, and hence increasing the density of the samples, requires further evaluations of the χ^2 function.

This method also suffers from poor scalability. As the number of model parameters increases, as does the dimensionality of the search space. Therefore, the number of χ^2 evaluations required increases exponentially with the number of model parameters. It would be desirable to have a method of minimisation that does not exhibit this loss of performance.

Given that each parameter of a reflection model lies in a continuous range of values, such

a discrete sampling of the parameter space to locate the optimal values would not appear to be practical. In searching for the optimal parameters it would be desirable for the search not to be limited by the discrete nature of the chosen method. For these reasons the brute force approach to locating the optimal model parameters is not considered further.

6.4.2 Gradient Search

By starting with an initial estimate of the model parameters, it is possible to traverse the χ^2 hypersurface such that a minimum is found. By considering the gradient of χ^2 with respect to each of the model parameters a path can be followed which will lead to a minimum. The method of *gradient search* [5] increments the parameters of an initial estimate such that χ^2 is reduced by following the path of steepest descent. The direction of maximum gradient, γ , is determined by evaluating the first order partial derivatives of the χ^2 function at each step,

$$\gamma_j = \left[-\frac{\partial \chi^2}{\partial a_j} \right] \delta a_j \quad j = 1, \dots, n, \quad (6.4)$$

where δa_j is an increment of the model parameter a_j .

This method of gradient search performs poorly as the search approaches the minimum. It is also the case that this method is easily misled by local minima in the hypersurface. Naively following the path of steepest descent does not ensure that the global minimum is located for any given starting estimate of the parameters. This can be overcome by repeated trials using different initial estimates of the parameters. Provided a suitable number of trials are performed, at suitably spaced starting points, the true global minimum should be located. The success of such a method is not assured, however.

6.4.3 Analytical Search

The method of gradient search performs iterative refinement of the initial parameters using a number of steps to find a minimum χ^2 . At each iteration the gradient must be evaluated with respect to each of the parameters. It would be desirable to perform a single step from the initial estimate and locate the minimum directly. This can be achieved by analysis of the χ^2 function, as given in Equation 6.2.

The χ^2 value of an initial set of parameters, \mathbf{a}_o , is denoted by χ_o^2 ,

$$\chi_o^2 = \sum_{i=1}^m \left[\frac{y_i - y(x_i; \mathbf{a}_o)}{\sigma_i} \right]^2. \quad (6.5)$$

The first order Taylor expansion of χ^2 about this point is given by

$$\chi^2 \approx \chi_o^2 + \sum_{j=1}^n \left(\frac{\partial \chi_o^2}{\partial a_j} \delta a_j \right), \quad (6.6)$$

where δa_j is an increment of the model parameter a_j . Equation 6.6 is at a minimum when the partial derivative with respect to each of the parameters is zero. Therefore, at the minimum, n equations are satisfied simultaneously,

$$\frac{\partial \chi^2}{\partial a_k} = \frac{\partial \chi_o^2}{\partial a_k} + \sum_{j=1}^n \left(\frac{\partial^2 \chi_o^2}{\partial a_j \partial a_k} \delta a_j \right) = 0 \quad k = 1, \dots, n. \quad (6.7)$$

Evaluation of the parameter increments, $\delta \mathbf{a}$, that satisfy Equation 6.7 will provide the optimal parameters, $\mathbf{a} + \delta \mathbf{a}$, which yield a minimum χ^2 .

The first and second order partial derivatives of χ_o^2 are obtained from Equation 6.5:

$$\frac{\partial \chi_o^2}{\partial a_k} = -2 \sum_{i=1}^m \frac{1}{\sigma_i^2} \left([y_i - y(x_i; \mathbf{a}_o)] \frac{\partial y(x_i; \mathbf{a}_o)}{\partial a_k} \right), \quad (6.8)$$

$$\frac{\partial^2 \chi_o^2}{\partial a_j \partial a_k} = 2 \sum_{i=1}^m \frac{1}{\sigma_i^2} \left(\frac{\partial y(x_i; \mathbf{a}_o)}{\partial a_k} \frac{\partial y(x_i; \mathbf{a}_o)}{\partial a_j} - [y_i - y(x_i; \mathbf{a}_o)] \frac{\partial^2 y(x_i; \mathbf{a}_o)}{\partial a_k \partial a_j} \right). \quad (6.9)$$

The set of linear simultaneous equations of Equation 6.7 can now be presented as a matrix equation,

$$\boldsymbol{\beta} = \delta \mathbf{a} \boldsymbol{\alpha}, \quad (6.10)$$

where,

$$\beta_k = -\frac{1}{2} \frac{\partial \chi_o^2}{\partial a_k}, \quad \alpha_{jk} = \frac{1}{2} \frac{\partial^2 \chi_o^2}{\partial a_j \partial a_k}. \quad (6.11)$$

The optimal parameter increments are given by solving the matrix equation,

$$\delta \mathbf{a} = \boldsymbol{\beta} [\boldsymbol{\alpha}]^{-1}. \quad (6.12)$$

The symmetric matrix of partial derivatives $\boldsymbol{\alpha}$ is termed the *curvature matrix* and represents a measure of the χ^2 surface curvature.

Once convergence has been achieved by this method a measure of the confidence in the result can be determined by inspection of the curvature matrix. The covariance of the fitted parameters is given by $[\boldsymbol{\alpha}]^{-1}$. Hence, the leading diagonal of this matrix gives a measure of confidence in each of the parameters [111].

This parabolic approximation of the χ^2 hypersurface is, in general, accurate close to the minimum. However, if the initial parameters, \mathbf{a}_o , are far from the optimal parameters the approximation fails. A common approach used to rectify this problem is to use the gradient search method far from the solution and resort to the parabolic approximation as the solution is approached. This combination of the two methods is provided by the Levenberg-Marquardt method [3, 83]. Many such Newton and quasi-Newton methods for non-linear regression exist. The Levenberg-Marquardt method is considered here as being representative of those various methods. For the purposes of this thesis, the implementation of the Levenberg-Marquardt method due to Press et al. [111] has been used.

Section 6.3 has shown that the χ^2 hypersurface for an image sequence contains local minima for certain illumination conditions and camera positions. It is therefore assumed that a search based on gradient descent will not be suitable for the minimisation of χ^2 for all image sequences. It will, however, provide an efficient method for the cases where the minimum is unique. The performance of this method, for models with more than two parameters, cannot be fully determined by considering the examples given in Section 6.3.

6.4.4 Simulated Annealing

First developed as a physical simulation of the cooling of crystalline structures, simulated annealing has become an established method for constraint satisfaction and combinatorial optimisation [138]. The algorithm, first developed by Metropolis et al. [86], provides a simulation of a collection of metallic atoms through stages of cooling. As the temperature is lowered the material structure attempts to achieve a state of thermal equilibrium such that the relative positions of the atoms in the crystalline lattice minimise the potential energy. The success with which an optimal, low energy, structure is reached is governed by the rate at which the material is cooled. By slow cooling near the material's freezing point the material is able to achieve its *ground state* of minimal energy. Rapid cooling around this critical temperature will cause non optimal crystals to form and the ground state will not be achieved. The resulting material may have defects and only locally optimal structures. The process by which materials are given time to achieve optimal low energy states as they cool is termed *annealing*. Metropolis' algorithm provides a statistically based simulation of the atoms' behaviour as cooling occurs.

Kirkpatrick et al. first recognised the similarity between statistical mechanics and combinatorial optimisation [63]. The search for a global minimum configuration to a constraint problem is analogous to that of obtaining the ground state by annealing. In this context the energy of the system is equated to the cost of a given combination of parameters. The simulated annealing algorithm for combinatorial problems requires an initial state of parameters, \mathbf{x}_o , the energy of which is evaluated by a *cost function*, E . The state is given a random perturbation, $\Delta\mathbf{x}$, and the resulting change in energy evaluated,

$$\Delta E = E(\mathbf{x} + \Delta\mathbf{x}) - E(\mathbf{x}). \quad (6.13)$$

If the new state has a lower energy it is accepted with probability 1, else it is accepted with a probability determined by the Boltzmann distribution,

$$p = ke^{(-\Delta E/T)}, \quad (6.14)$$

where T is the current temperature and k is the Boltzmann constant. As the temperature is decreased the probability of changing to a higher energy state is reduced. Provided the system is allowed to cool sufficiently slowly the optimal, minimal cost, parameter combination should be achieved. The efficiency of the algorithm in finding the global minimum is penalised by the increased number of cost function evaluations required over that associated with gradient based optimisation methods. The technique does, however, have the ability to ignore local topology at high temperatures to find the region of low cost. As temperature is reduced the solution is further refined.

Though originally designed for the optimisation of problems where each parameter may take on a discrete value, the method of simulated annealing has since been extended to problems defined in a continuous domain [139]. Here, the annealing method can be used to locate the global minimum of a function of many variables. The state of the system is analogous to a position on the function hypersurface. Randomly perturbing this state provides a random walk about the domain of the function. Using the Boltzmann distribution, the state can be progressed out of local minima such that the global minimum is located. The method has been successfully employed for a number of optimisation tasks [8, 15, 139]. It has also been shown to provide a regression technique for non-linear least-squares fitting problems [139].

Press et al. [111] present a variation on traditional simulated annealing which is cited as more efficient than the methods given above. They propose the use of a geometric simplex of

points to represent the system state. This simplex lies on the hypersurface of the function to be minimised, in this case the χ^2 function. As such, the simplex has $n + 1$ vertices where n is the number of parameters. The simplex is allowed to undergo a number of transformations. These are depicted in Figure 6.23. In each case the highest vertex of the simplex is translated toward the lowest face. The magnitude of the translation is determined by the relative decrease in cost thus achieved. At zero temperature the simplex is allowed to move such that it traverses the hypersurface following the path of steepest descent. At higher temperatures the simplex is able to accept translations which result in a higher cost. This is achieved by perturbing the simplex vertices in relation to temperature and thus providing the simplex with a Brownian motion which allows it to escape local minima. At each temperature the simplex is allowed to undergo a number of moves to find a low energy state.




Figure 6.23: *Simplex translations: (a) original position of the simplex, vertex 1 has the highest cost; (b) reflection; (c) reflection and expansion; (d) contraction.*

The efficiency of any simulated annealing approach is governed by correct selection of a cooling schedule. This determines the initial start temperature, T_0 , the rate at which this temperature is reduced and the stop criterion for halting the search. In the case of the simplex approach it also necessary to define how many translations the simplex may undergo at any given temperature. Many cooling schedules have been proposed and the performance of each is problem specific [138]. This work uses the simple schedule proposed by Kirkpatrick et al. [62]. Here, the initial temperature is reduced by a constant factor α after every n translations of the simplex,

$$T_{i+1} = \alpha T_i \quad i = 0, 1, 2, \dots, \quad (6.15)$$

where α is a constant smaller than, but close to, one. The selection of T_0 is again problem specific. It should be high enough such that all states are reachable from the initial position of the simplex, but not excessively high so as to cause unnecessary computations [138]. Ideally the system should be cooled slowly in the temperature region where the energy of the system decreases most rapidly. This is called the *phase transition* and is analogous to the freezing point of a material. Figure 6.24 shows the expected change of energy as a system is cooled.




Figure 6.24: Graph of average system energy plotted against temperature, with phase transition region indicated.

It is proposed that the method of simulated annealing is suitable for all the χ^2 examples given in Section 6.3 since it can escape the local minima that have been shown to exist in these cases. Further to this, it should be equally applicable to models with more than two parameters. Whilst the Levenberg-Marquardt method may prove to be efficient, it is expected that it will not be suitable for all image sequences.

6.4.5 Limiting the Search Space

As stated in Section 3.4.5, each of the reflection model parameters has a finite range. It has been found that the technique used to constrain each of the above search methods, so that resulting estimated parameters are kept within these ranges, has significant affect on performance. Bounding the permitted parameter space by a plateau of high χ^2 value presents difficulties when the optimal parameters lie on or near to a permitted range limit. In this case the minimum presented at the optimal parameters may not be adequately approximated by a parabolic, as required by the Levenberg-Marquardt method. Difficulties with such a bounding technique have also been experienced in the use of simulated annealing. Again, it would appear that the possibility of high χ^2 function values surrounding the region of a minimum compromises the ability of the simplex to converge satisfactorily.

It has been found that the regression techniques exhibit improved performance if the permitted parameter space is bounded by a monotonically increasing χ^2 . In the case of the Levenberg-Marquardt regression this can be shown to be more suitable for the parabolic surface assumption. In the case of simulated annealing, the simplex is able to move beyond the permitted parameter

values but will descend the gradient towards the minimum and shall therefore result in parameters which are within the permitted values.

6.5 Summary

This chapter has shown that a sequence of images, obtained from a static camera over a period of time, may be used to determine the parameters of a reflection model which best characterise the observed material. It is expected that the temporal change in image pixel values, due to changing natural illumination, will provide suitable information for such parameters to be derived. The limitations of this assumption have been assessed and it has been shown that, in some cases, an accurate estimate may not be possible using the proposed methods.

This chapter has presented a measure of the difference between the temporal change in image pixel values and that predicted by a model. This measure, χ^2 , allows for the comparison of image data with a model of the processes by which light reflected from a scene creates an image. The minimisation of the difference between such measured values and those predicted by a model allows a best fit model to be estimated. In this context, surface reflection is measured using an image based system. Therefore, the data is comprised of image pixel values. As such, the model used to describe the pixel values calls upon models of camera response and natural illumination in combination with models of surface reflection. With the parameters of these extra models known *a priori*, the free parameters of the combined model are those of that used to describe the reflection characteristics of the observed surface. Regression techniques allow the estimation of these free parameters and hence the surface reflection model for the observed material may be determined.

Based upon analysis of the χ^2 function for a variety of image sequences two regression methods have been proposed. These methods should allow for the estimation of best fit parameters, in a least-squares sense, to the observed data. The performance of these two methods shall be compared for an increasing number of free parameters in the results.

Chapter 7

Results

This chapter presents results obtained using two regression methods with synthetic as well as real image sequences. The use of synthetic image sequences allows the performance of each of the techniques to be understood in the context of clean data obtained in known conditions. Work on synthetic sequences cannot, however, replace observations made on real data obtained with a camera. For this reason an image sequence obtained from a clear day and consisting of a number of material types is considered.

The synthetic image sequences have been produced using the *Radiance* lighting simulation system. This has been described in Chapters 4 and 6 and allows for an accurate simulation of illumination by sun and sky sources. The results presented here focus on the analysis of scenes illuminated by a clear sky. The illumination therefore corresponds to a visible sun and a CIE clear sky, as presented in Chapter 4. The illumination models used to generate the synthetic image sequences are the same as those used in the regression techniques. Also, these results only consider the analysis of a single pixel in each image of a sequence. Further to this, only horizontal planar surfaces are analysed in each case. The effects of alternative illumination and surface orientations have been described in Chapter 6.

Actual data obtained by a camera has been obtained from a sequence of images viewing a collection of planar horizontal surfaces. These material samples have been illuminated by a natural clear sky. Again, only a single pixel from the projection of each sample is considered.

This chapter focuses on the use of Ward's reflection model. Whilst the techniques developed

in Chapter 6 for the estimation of model parameters are not linked to a particular model, the Ward model is chosen here for its applicability to a wide variety of surface types. Three variants of the basic Ward model, each with an increasing number of parameters, are considered, and these are summarised in Section 7.1.

7.1 Overview of the Reflection Model

The reflection model considered in this chapter is that due to Ward [141], as described in Section 3.4.5. It is briefly summarised here for convenience. The bidirectional reflectance-distribution function, f_r , is modelled by,

$$f_r(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{\rho_d}{\pi} + \rho_s \cdot \frac{1}{\sqrt{\cos \theta_i \cos \theta_r}} \cdot \frac{\exp(-\tan^2 \delta / \sigma_\alpha^2)}{4\pi\sigma_\alpha^2}, \quad (7.1)$$

where ρ_d is the normal lobe reflectance, ρ_s is the foreshatter lobe reflectance, σ_α is the rms roughness of the surface, and δ is the *half angle*. This chapter considers three variants of this basic model, each with an increasing number of parameters.

A two-parameter model requires that $\rho_d + \rho_s = 1$ and as such the total reflectance of surfaces modelled by this function are constant. The resulting parameters to be estimated for this model are, therefore, ρ_s and σ_α .

A three-parameter achromatic model allows for the inclusion of a total reflectance term, R . This allows control over the total reflectance of a surface as well as the proportion of reflection into each of the reflection lobes. Again, $\rho_d + \rho_s = 1$, and the three parameters to be estimated are R , ρ_s and σ_α .

A five-parameter model replaces the total reflectance term, R , with a material colour vector, \mathbf{C} . This vector gives the reflectance of the surface at three discrete wavelengths, \mathbf{C}_{red} , \mathbf{C}_{green} and \mathbf{C}_{blue} . How this colour vector is used in the model depends upon the value of ρ_s . It allows for the modelling of coloured reflection from plastic and metallic surfaces. The five parameters to be estimated for this model are, therefore, \mathbf{C}_{red} , \mathbf{C}_{green} , \mathbf{C}_{blue} , ρ_s and σ_α . The permitted range for σ_α is (0.0 – 0.2) whilst all other parameters have a range of (0.0 – 1.0).

7.2 Levenberg-Marquardt Regression

The Levenberg-Marquardt method of regression (see Section 6.4.3) provides a gradient descent based approach of the χ^2 surface from an initial estimate of the parameters. It has been shown in Chapter 6 that this regression method may be misled by local minima in the χ^2 surface; however, in the two-parameter model these local minima were not apparent in all cases. This section will look at the performance of this technique when considering the two and three-parameter reflection models.

7.2.1 Two-Parameter Reflection Model

This section considers the performance of the Levenberg-Marquardt technique on three synthetic test sequences. These are selected from those examined in Chapter 6. Two of these examples have been shown not to contain non-optimal minima and have gradients which descend toward the location of the optimal parameters. As such, it would be expected that this method would perform well in these cases. A third example, shown to exhibit local minima, is also considered.

Table 7.1 shows the performance of this technique when used on the sequence observing a Lambertian disc from the north. The χ^2 values for this example are shown in Figure 6.5. Given that the global minimum lies at the optimal parameters, and that this minimum is unique in this case, it would be expected that this gradient descent method would perform well. Two observations can be made from these results. There are a number of cases where the method has failed to converge upon a solution and the method has not consistently found the optimal parameters.

Failure to converge, in this case, is due to the curvature matrix becoming singular, hence the matrix α of Equation 6.12 may not be inverted to find the necessary parameter increments $\delta\mathbf{a}$. This has occurred owing to the nature of the Ward reflection model near $\rho_s = 0$. In the case of the Ward model, when $\rho_s = 0$ the roughness parameter, σ_α , has no significance. It is therefore the case that, as the method converges towards $\rho_s = 0$, the partial derivative $\frac{\partial\chi^2}{\partial\sigma_\alpha} \rightarrow 0$. This leads to an asymmetric curvature matrix and the method fails. This highlights an inherent problem with using regression methods that rely on the evaluation of partial derivatives. If the data is such that a change in value of a model parameter has no effect on the value of χ^2 , the method is prone to failure since it cannot determine the direction of steepest descent.

The results in Table 7.1 show that, as the distance of the initial, starting estimate from the optimal parameters is increased, the performance of the method degenerates. This behaviour is clearly shown in Figure 7.1. In this graph the error, based on distance in the parameter space from the solution, of the initial estimate is compared with the error of the final result. Points below the dashed diagonal indicate an improvement of the initial estimate toward the expected result. It can be seen that, for initial estimates which are close to the optimal parameters, the method converges to the expected values. However, as the distance of the initial estimate from the optimal parameters is increased the error in the estimated parameters also increases. Such dependence upon the initial parameter estimates is a characteristic of this method [5], though the extent to which it affects the result is problem specific.

Table 7.1: *Levenberg-Marquardt results for two-parameter model. Optimal parameter values are $\rho_s = 0.0$ where σ_α has no significance.*

Initial		Result		Variance		χ^2 of	χ^2
ρ_s	σ_α	ρ_s	σ_α	ρ_s	σ_α	result	evaluations
0.1	0.00	0.04	0.00	0.095	0.101	11	65
0.0	0.10	Fail		-	-	-	-
0.1	0.10	Fail		-	-	-	-
0.2	0.05	Fail		-	-	-	-
0.2	0.20	0.01	0.20	0.007	0.063	41	39
0.4	0.05	0.00	0.00	0.358	0.126	11	52
0.4	0.15	Fail		-	-	-	-
0.6	0.05	0.00	0.00	0.233	0.806	11	52
0.6	0.20	Fail		-	-	-	-
0.8	0.00	0.07	0.00	0.090	0.077	1527	52
0.8	0.15	0.00	0.19	0.023	1.362	11	65
1.0	0.10	Fail		-	-	-	-
1.0	0.20	0.00	0.20	35.147	47.925	11	65

Table 7.2 considers a rough metallic disc observed from the north. The χ^2 values for this example are shown in Figure 6.6(c). Again, as with the above example, the surface descends toward the optimal parameters of $\rho_s = 1$, $\sigma_\alpha = 0.2$. In this case failure to converge, owing to a singular curvature matrix, has only occurred when the initial estimate of the parameters is $\rho_s = 0$. However, it can again be seen that, as the distance of the initial estimate from the optimal parameters is increased, the performance of the method degenerates. This behaviour can be clearly seen in Figure 7.2.




Figure 7.1: Error in resulting parameters against error of initial estimate for two-parameter model. Sequence observes a Lambertian surface from the north.

Table 7.2: Levenberg-Marquardt results for two-parameter model. Optimal parameter values are $\rho_s = 1.0$ and $\sigma_\alpha = 0.2$.

Initial		Result		Variance		χ^2 of result	χ^2 evaluations
ρ_s	σ_α	ρ_s	σ_α	ρ_s	σ_α		
1.0	0.15	1.00	0.20	0.023	0.002	15	52
0.9	0.20	1.00	0.20	25.219	119.174	15	39
0.9	0.10	1.00	0.20	22.188	11.951	15	52
0.8	0.15	1.00	0.19	0.059	0.000	17	39
0.8	0.05	0.80	0.05	0.001	0.009	103348	39
0.6	0.10	1.00	0.20	137.527	61.821	15	52
0.6	0.00	0.60	0.00	0.002	0.026	130216	39
0.4	0.15	1.00	0.20	19.474	78.482	15.23	39
0.4	0.05	0.40	0.05	0.002	0.027	200227	39
0.2	0.20	1.00	0.00	13.960	207.219	148269	52
0.2	0.05	0.20	0.05	0.003	0.062	289681	39
0.0	0.05	Fail		-	-	-	-
0.0	0.15	Fail		-	-	-	-




Figure 7.2: Error in resulting parameters against error of initial estimate for two-parameter model. Sequence observes a rough metallic surface from the north.

Table 7.3 considers the case of a smooth metallic disc observed from the east. Figure 6.7(b) showed that, in this case, there exists a non-optimal minimum in the surface of the χ^2 function. The problems that this presents are shown in the results for this experiment. It can be seen that, for initial starting positions that are contained within the region of this local minimum, the gradient descent method is unable to locate the global minimum. This can be seen with many of the poor results being in the region of $\rho_s = 0.2, \sigma_\alpha = 0.2$. As with the above example, where the optimal parameters are reachable using a gradient descent method, the accuracy of the method still depends upon the distance of the initial estimate from the optimal parameter values. Figure 7.3 clearly shows this behaviour. This figure also shows a cluster of results corresponding to the non-optimal minimum.

Since the Levenberg-Marquardt method provides an estimate of the confidence in the resulting parameters it would be hoped that, in the cases where the result has a high error, this would be indicated by the values of the curvature matrix on convergence. However, as Figure 7.4 shows, there is no observable correlation between the confidence in the resulting parameters indicated by the curvature matrix and the known error in the result. This is partly to be expected since the confidence values only give an indication of how well the hypersurface at the minimum matches the parabolic assumption. A non-optimal minimum in the hypersurface may be more adequately approximated by a parabolic than that of the optimal global minimum.

Table 7.3: Levenberg-Marquardt results for two-parameter model. Optimal parameter values are $\rho_s = 1.0$ and $\sigma_\alpha = 0.0$.

Initial		Result		Variance		χ^2 of	χ^2
ρ_s	σ_α	ρ_s	σ_α	ρ_s	σ_α	result	evaluations
1.0	0.05	1.00	0.02	0.069	0.002	116	52
0.9	0.00	1.00	0.00	0.064	0.005	8	39
0.9	0.10	1.00	0.08	0.019	0.001	203974	39
0.8	0.15	0.45	0.20	0.006	0.017	140935	39
0.8	0.05	1.00	0.02	0.038	0.000	116	52
0.6	0.10	0.60	0.10	0.005	0.000	34596	39
0.6	0.00	1.00	0.00	0.000	0.000	8	52
0.4	0.15	0.28	0.20	0.004	0.021	127466	39
0.4	0.05	1.00	0.00	0.004	0.006	8	39
0.2	0.20	0.20	0.20	0.006	0.007	122432	39
0.2	0.05	0.47	0.00	0.018	0.166	50783	39
0.0	0.05	0.37	0.00	0.022	0.052	63979	39
0.0	0.15	0.10	0.15	0.004	0.004	123292	39




Figure 7.3: Error in resulting parameters against error of initial estimate for two-parameter model. Sequence observes a smooth metallic surface from the east.




Figure 7.4: Variance of resulting parameter estimates. Note log scale on ordinate axis.

7.2.2 Achromatic Reflection Model

This section considers the three-parameter achromatic reflection model and synthetic greyscale image sequences. The topology of the χ^2 surface for such cases has not been considered in Chapter 6, though it can be assumed that it is at least as complex as observations from the two-parameter model suggested.

Figure 7.5 shows results obtained from a sequence observing a Lambertian disc viewed from the east. The optimal parameters for this sequence are $\rho_d = 0.33$, $\rho_s = 0.0$ and $\sigma_\alpha = 0.05$. As with the two-parameter results given in Section 7.2.1, it can be seen that the performance of this method is dependent upon the accuracy of the initial parameter estimates. This trend is repeated for the examples shown in Figure 7.6 and 7.7. In these examples the number of χ^2 evaluations required for convergence were similar to those of the two-parameter examples above.

7.2.3 Summary of Levenberg-Marquardt Performance

The number of χ^2 evaluations required for convergence in these examples is low compared to that which would be required by a brute force search of the parameter space. To sample the value of χ^2 at ten per cent increments in each dimension would require 10^2 and 10^3 samples with the two and three-parameter reflection models respectively. The Levenberg-Marquardt method has required significantly fewer evaluations to obtain a result. However, the quality of this result has been shown to be dependent upon the initial estimate of the parameter values. Given that such initial




Figure 7.5: Error in resulting parameters against error of initial estimate for three-parameter model. Sequence observes a plastic surface with parameters $R = 0.33$, $\rho_s = 0.2$, $\sigma_\alpha = 0.05$ from the east.




Figure 7.6: Error in resulting parameters against error of initial estimate for three-parameter model. Sequence observes a metallic surface with parameters $R = 0.66$, $\rho_s = 0.8$, $\sigma_\alpha = 0.15$ from the east.




Figure 7.7: Error in resulting parameters against error of initial estimate for three-parameter model. Sequence observes a metallic surface with parameters $R = 1.0$, $\rho_s = 1.0$, $\sigma_\alpha = 0.2$ from the east.

estimates of the parameters are not available to the system, the method proves to be inappropriate for this application. This could be overcome by performing several repeated applications of the method using uniformly distributed starting positions. This would require an increased number of χ^2 evaluations and, depending on the density of the starting positions, may not find the optimal parameters.

It has also been shown that, in some cases, the method is unable to converge upon a solution. Whilst this is due to the nature of the Ward model, similar conditions may also apply to alternative reflection models. Given the poor performance of the Levenberg-Marquardt regression method in these experiments it has not been used in the case of the five-parameter colour reflection model, nor has it been used on the real image sequences.

7.3 Simulated Annealing

This section considers the performance of simulated annealing for regression. The Levenberg-Marquardt method has shown three distinct problems in this application: the dependence on the initial parameter estimates; the possibility of a singular curvature matrix; and the problem of non-optimal minima. It is expected that simulated annealing, described in Section 6.4.4, will overcome these three issues with respect to this application.

Since this method does not rely on the calculation of partial derivatives the situation where

$\rho_s = 0$ does not present the problem it did for the Levenberg-Marquardt method. If the simplex does enter a plateau region, the thermal motion of the simplex will allow it to traverse the χ^2 surface.

The simulated annealing method will randomly search the entire parameter space. Therefore, the accuracy of the solution should not be dependent upon the initial parameter estimates. This will be true provided that the selected cooling schedule allows the simplex to escape local minima and to traverse the χ^2 surface such that the global minimum is located. As stated in Section 6.4.4, the method of simulated annealing is able to ignore local topology at high temperatures and is, therefore, able to locate the optimal parameters which yield a minimum χ^2 value.

7.3.1 Selection of Cooling Schedule

The performance of simulated annealing is highly dependent on the selection of an appropriate cooling schedule. This section looks at the selection of a cooling schedule suitable for the two-parameter reflection model considered in the examples of Chapter 6. As stated in Section 6.4.4, the cooling schedule should be such that, at the initial temperature, T_0 , all points on the surface are reachable with equal probability.

To understand the significance that the cooling schedule has on the performance of this method, the case of a smooth metallic disc observed from the east is considered. The χ^2 function values for this example are shown in Figure 6.7(b). In each of the experiments that follow, the initial position of the simplex is in the local, non-optimal, minimum indicated in Figure 6.7(b). Figure 7.8 shows the average energy of the simplex as the system is cooled from three different values of T_0 . It can be seen that, if the initial temperature is too low, as shown in Figure 7.8(a), the simplex is not provided with sufficient energy with which to escape the local minimum in which it starts. In this example the optimal parameters are not correctly estimated by the method. Starting the system with a high temperature, as in Figure 7.8(c) does allow the simplex to escape the local minimum and the global minimum is located as the system is cooled. However, the average energy of the simplex remains high and rapidly reduces towards the end of the cooling. This is due to the simplex having an initially high level of energy and exhibiting erratic behaviour regardless of the gradient of the χ^2 surface. At these high temperatures the simplex is randomly sampling points on the χ^2 surface and is not attempting to refine the solution. Only when the temperature

has sufficiently reduced is the simplex able to descend toward the global minimum. Whilst such a cooling schedule does enable the global minimum to be located, it is at the expense of increased evaluations of χ^2 . Figure 7.8(b) shows an improved cooling schedule. Here the simplex is started at a sufficiently high temperature to escape the local minimum but is reduced sufficiently slowly so as to refine the solution.




Figure 7.8: Effect of initial temperature on simulated annealing performance: (a) $T_0 = 10^5$; (b) $T_0 = 10^6$; (c) $T_0 = 10^7$.

The rate of cooling from the initial temperature should allow sufficient time in the phase transition region shown in Figure 6.24. This is controlled by two parameters of the cooling schedule: the rate at which the system is cooled, α , and the number of iterations, m , that are allowed at each temperature. As stated in Section 6.4.4, the temperature is reduced by a factor α at each stage of the cooling. Figure 7.8(b) shows cooling with $\alpha = 0.7$ and $m = 10$. Alternative values for these two parameters are shown in Figure 7.9.

Figure 7.9(a) shows cooling whereby increased iterations are allowed at each temperature. This has shown no improvement in the accuracy of the result and has come at an increased number of evaluations of χ^2 . Figure 7.9(b) shows cooling with $m = 5$. In this case the system has not been allowed sufficient iterations with which to attempt to find an equilibrium at each temperature. The inability to refine the solution has resulted in the final simplex position not being at the global minimum. This can again be seen in Figure 7.9(c), however, this has been the result of cooling too rapidly with $\alpha = 0.6$. Here, the system has been quenched and has not been able to settle into the optimal state. Figure 7.9(d) shows cooling with $\alpha = 0.8$. Here the temperature is reduced slowly but the solution is not improved over that obtained with $\alpha = 0.7$. The consequence of using a reduced cooling rate has been the increased number of χ^2 evaluations required.

This section has shown the importance of the selection of an appropriate cooling schedule in the performance of the method. Incorrect selection of the various parameters may lead to non




Figure 7.9: Effect of cooling rate on simulated annealing performance: (a) $\alpha = 0.7$, $m = 20$; (b) $\alpha = 0.7$, $m = 5$; (c) $\alpha = 0.8$, $m = 10$; (d) $\alpha = 0.6$, $m = 10$.

optimal solutions or to unnecessary evaluations of the χ^2 function. Where a sequence contains a large number of images, or a number of pixels are considered within each image, the repeated evaluations of χ^2 may prove to be computationally expensive.

7.3.2 Two-Parameter Reflection Model

The performance of simulated annealing for the two-parameter model and synthetic image sequences is now considered. In each of the examples presented here, the cooling schedule used is $T_0 = 10^6$, $\alpha = 0.7$, $m = 10$. Section 7.3.1 has shown this to be a suitable cooling schedule for this model.

Table 7.4 shows results for a Lambertian disc viewed from the north. These results should be compared with those presented in Table 7.1. It can be seen that the optimal parameters have been correctly estimated by this method, and that the accuracy of the result has not been determined by the initial estimate of the parameters. Note that the value of σ_α is not significant when $\rho_s = 0$, as it does in this example. It can be seen, however, that using this cooling schedule has required a significantly increased number of χ^2 evaluations over that required by the Levenberg-Marquardt method.

Table 7.4: *Simulated annealing results for two-parameter model.*
Optimal parameter values are $\rho_s = 0.0$ where σ_α has no significance.

Initial		Result		χ^2 evaluations
ρ_s	σ_α	ρ_s	σ_α	
0.90	0.18	0.00	0.19	464
0.58	0.02	0.00	0.14	463
0.78	0.00	0.00	0.00	462
0.16	0.02	0.00	0.04	461
0.62	0.01	0.00	0.07	434
0.34	0.16	0.00	0.11	463

Table 7.5 and Table 7.6 show results for differing materials. These results should be compared to Tables 7.2 and 7.3 respectively which consider the same image sequences. Again, it can be seen that the accuracy of the result is not dependent upon the initial estimate. It can also be seen that the method has not been misled by local, non-optimal, minima. These results show little variance in the estimated parameters.

Table 7.5: *Simulated annealing results for two-parameter model.*
Optimal parameter values are $\rho_s = 1.0$ and $\sigma_\alpha = 0.2$.

Initial		Result		χ^2 evaluations
ρ_s	σ_α	ρ_s	σ_α	
0.40	0.04	1.0	0.20	311
0.33	0.04	1.0	0.20	376
0.31	0.10	1.0	0.20	409
0.02	0.07	1.0	0.20	390
0.86	0.01	1.0	0.19	364
0.65	0.18	1.0	0.20	367

Figure 7.10 shows a typical random walk for the smooth metallic disc sequence considered for Table 7.6. It can be seen that, despite being started in a local minimum, the nature of the annealing process has allowed the simplex to escape and locate the global minimum.

7.3.3 Achromatic Reflection Model

This section considers the performance of the simulated annealing method on the achromatic reflection model and synthetic greyscale image sequences. This reflection model has three param-

Table 7.6: Simulated annealing results for two-parameter model. Optimal parameter values are $\rho_s = 1.0$ and $\sigma_\alpha = 0.0$.

Initial		Result		χ^2 evaluations
ρ_s	σ_α	ρ_s	σ_α	
0.03	0.01	1.0	0.0	354
0.78	0.18	1.0	0.01	334
0.31	0.10	1.0	0.0	321
0.23	0.07	1.0	0.0	401
0.87	0.12	1.0	0.0	364
0.11	0.05	1.0	0.0	334




Figure 7.10: Typical random walk over a χ^2 surface.

eters and therefore requires the simplex to traverse a higher dimension space as that considered above. Owing to this higher dimensionality, it has been found that slower cooling has provided consistent results for this reflection model. For these examples the selected cooling schedule has been $T_0 = 10^6$, $\alpha = 0.8$ and $m = 10$.

Table 7.7 shows results for a plastic disc with Ward reflection model parameters $R = 0.33$, $\rho_s = 0.20$ and $\sigma_\alpha = 0.05$. It can be seen that, in each case, the method has correctly estimated the optimal parameters regardless of the initial estimate of the parameters. These results should be compared to those shown in Figure 7.5 for the Levenberg-Marquardt method on the same image sequence.

Table 7.7: *Simulated annealing results for three-parameter model.*
Optimal parameter values are $R = 0.33$, $\rho_s = 0.2$ and $\sigma_\alpha = 0.05$.

Initial			Result			χ^2 evaluations
R	ρ_s	σ_α	R	ρ_s	σ_α	
0.64	0.22	0.05	0.33	0.21	0.05	719
0.77	0.49	0.14	0.33	0.20	0.05	725
0.50	0.08	0.16	0.33	0.20	0.05	717
0.60	0.22	0.17	0.33	0.20	0.05	719
0.85	0.80	0.03	0.33	0.20	0.05	721
0.11	0.20	0.14	0.33	0.20	0.05	728

Table 7.8 shows results for a metallic disc with Ward model parameters $R = 0.66$, $\rho_s = 0.80$ and $\sigma_\alpha = 0.15$. It can be seen that, in each case, the method has correctly estimated the optimal parameters. Again, these results should be compared to the performance of the gradient descent method on this image sequence, shown in Figure 7.6.

Table 7.8: *Simulated annealing results for three-parameter model.*
Optimal parameter values are $R = 0.66$, $\rho_s = 0.8$ and $\sigma_\alpha = 0.15$.

Initial			Result			χ^2 evaluations
R	ρ_s	σ_α	R	ρ_s	σ_α	
0.40	0.15	0.20	0.66	0.80	0.15	718
0.48	0.80	0.14	0.66	0.80	0.15	715
0.32	0.03	0.09	0.66	0.80	0.15	711
0.88	0.85	0.10	0.66	0.80	0.15	712
0.18	0.78	0.06	0.66	0.81	0.15	719
0.09	0.48	0.17	0.67	0.80	0.15	719

Table 7.9 shows results for a metallic disc with Ward model parameters $R = 1.0$, $\rho_s = 1.0$ and $\sigma_\alpha = 0.0$. Again, it can be seen that, in each case, the method has correctly estimated the optimal parameters. These results should be compared to those shown in Figure 7.7 for the same image sequence.

Table 7.9: *Simulated annealing results for three-parameter model. Optimal parameter values are $R = 1.0$, $\rho_s = 1.0$ and $\sigma_\alpha = 0.0$.*

Initial			Result			χ^2 evaluations
R	ρ_s	σ_α	R	ρ_s	σ_α	
0.54	0.01	0.06	1.0	1.0	0.0	725
0.74	0.63	0.08	0.99	1.0	0.0	736
0.77	0.72	0.11	1.0	1.0	0.0	730
0.28	0.84	0.01	1.0	0.99	0.0	718
0.42	0.78	0.16	1.0	1.0	0.01	727
0.87	0.24	0.13	1.0	1.0	0.0	730

These results have shown significant improvement over the gradient descent method considered in Section 7.2.1. As with the two-parameter model, these results show little variation in the final estimated parameter values.

7.3.4 Colour Reflection Model

The performance of simulated annealing with colour images and the five-parameter colour reflection model is now considered. In this five dimensional parameter space it has been found that a cooling schedule of $T_0 = 10^7$, $\alpha = 0.8$ and $m = 20$ provides consistent results. Each of the image sequences considered here has been synthetically generated. Results for three such sequences are presented.

Tables 7.10, 7.11 and 7.12 show results for horizontal plastic surfaces with differing reflection model parameters. It can be seen that, in each case, the method has correctly estimated the optimal parameter values. As with the previous reflection model examples, there is little variation in the estimated parameter values in each case.

Table 7.10: *Simulated annealing results for five-parameter model.*
 Optimal parameter values are $C_{red} = 0.25$, $C_{green} = 0.50$, $C_{blue} = 0.75$,
 $\rho_s = 0.2$ and $\sigma_\alpha = 0.05$.

Initial					Result					χ^2
C_{red}	C_{green}	C_{blue}	ρ_s	σ_α	C_{red}	C_{green}	C_{blue}	ρ_s	σ_α	evaluations
0.67	0.48	0.47	0.09	0.12	0.25	0.50	0.75	0.20	0.05	1561
0.99	0.94	0.64	0.31	0.11	0.25	0.50	0.74	0.20	0.05	1574
0.13	0.78	0.18	0.23	0.03	0.25	0.50	0.74	0.20	0.05	1574
0.30	0.30	0.18	0.15	0.16	0.25	0.50	0.75	0.21	0.05	1405
0.85	0.65	0.61	0.18	0.03	0.25	0.50	0.75	0.20	0.05	1561
0.65	0.51	0.73	0.87	0.01	0.24	0.50	0.75	0.21	0.05	1574

Table 7.11: *Simulated annealing results for five-parameter model.*
 Optimal parameter values are $C_{red} = 0.75$, $C_{green} = 0.50$, $C_{blue} = 0.25$,
 $\rho_s = 0.8$ and $\sigma_\alpha = 0.15$.

Initial					Result					χ^2
C_{red}	C_{green}	C_{blue}	ρ_s	σ_α	C_{red}	C_{green}	C_{blue}	ρ_s	σ_α	evaluations
0.92	0.57	0.87	0.20	0.19	0.75	0.50	0.25	0.81	0.15	1401
0.69	0.31	0.95	0.53	0.05	0.75	0.49	0.24	0.80	0.15	1575
0.88	0.85	0.52	0.40	0.05	0.75	0.50	0.25	0.81	0.15	1478
0.11	0.64	0.77	0.70	0.15	0.75	0.50	0.25	0.81	0.15	1575
0.91	0.89	0.16	0.26	0.19	0.74	0.49	0.25	0.82	0.15	1575
0.09	0.38	0.86	0.94	0.19	0.74	0.49	0.24	0.81	0.15	1575

Table 7.12: *Simulated annealing results for five-parameter model.*
 Optimal parameter values are $C_{red} = 1.0$, $C_{green} = 1.0$, $C_{blue} = 1.0$,
 $\rho_s = 1.0$ and $\sigma_\alpha = 0.0$.

Initial					Result					χ^2
C_{red}	C_{green}	C_{blue}	ρ_s	σ_α	C_{red}	C_{green}	C_{blue}	ρ_s	σ_α	evaluations
0.15	0.54	0.08	0.17	0.03	0.99	1.0	1.0	1.0	0.0	1564
0.01	0.21	1.0	0.70	0.10	0.99	1.0	1.0	1.0	0.0	1574
0.99	0.28	0.0	0.24	0.14	1.0	0.98	1.0	1.0	0.0	1575
0.46	0.59	0.32	0.29	0.02	0.98	0.98	1.0	1.0	0.0	1574
0.66	0.68	0.84	0.75	0.10	0.99	1.0	0.99	1.0	0.0	1575
0.73	0.62	0.77	0.33	0.13	1.0	1.0	1.0	1.0	0.0	1574

7.3.5 Summary of Simulated Annealing Performance

This section has demonstrated the effectiveness of simulated annealing for the estimation of reflection model parameters. The results are much improved over those obtained using the gradient descent approach shown in Section 7.2 and do not exhibit the problems encountered with that method.

The quality of these results is to be expected, however. These trials have been performed using synthetic image sequences, observing surfaces generated using the same reflection model. It is also the case that, in each example, the camera and illumination models used are the same. There is, therefore, no reason why, given a suitable cooling schedule, this method shall fail to estimate the optimal parameter values in each of the above cases. The main purpose of these experiments has been to validate the regression technique in this application.

The improved results obtained by simulated annealing have come at the expense of an increased number of χ^2 function evaluations over that required by the Levenberg-Marquardt method. However, the number of evaluations should be compared with that required by a brute force search of the parameter space. In the case of the five-parameter colour reflection model, a brute force search would require $10^5 \chi^2$ function evaluations to sample the parameters at increments of 10 per cent in each dimension. This can be compared with the average of 1548 function evaluations required by simulated annealing. In the case of simulated annealing with the five-parameter model, a typical execution time is in the order of four hours¹.

7.4 Real Image Sequence

This section considers real data obtained using a photographic camera, with images subsequently scanned to provide a digital image. Whilst the camera model developed in Chapter 5 is not directly applicable to these images, it has been found that the images obtained by this method are suitable for analysis. Calibration of these images has been achieved through the use of an ANSI IT8 colour target within each image. This has been used to ensure linearity in response and colour reproduction across the image sequence.

¹Timings obtained on SiliconGraphics Indy.

The image sequence consists of thirteen images, taken half-hourly. The scene was observed on August 16 with the first image taken at 12:00 BST (British Summer Time) and the final image at 18:00 BST. The location of the scene was 53°N, 1°W and with sky conditions described as a clear sky with little observable pollution. The camera is located in the east and looks down upon the sample area at an angle of 37° to the surface normal. The scene is composed of eight material samples, horizontally orientated, to which the full sky hemisphere is visible. A typical image from the sequence is shown in Figure 7.11. The eight material samples observed in this sequence are described in Table 7.13.



Figure 7.11: An example image from a sequence observing a selection of sample materials on a clear day. The camera is situated in the east.

Table 7.13: Description of sample materials.

1	White paper.
2	Pale blue paper.
3	White gloss paint.
4	Polished mild steel.
5	Mild steel roughened using a 400 grit abrasive.
6	Mild Steel roughened using a 40 grit abrasive.
7	Yellow gloss paint.
8	Red gloss paint.

Figure 7.12 shows the change in colour pixel value, for a single pixel within each of the material samples, across the image sequence. It can be seen that the data contains noise but a difference in each of the graphs can be observed. It should also be noted that there is an increased blue pixel response. This is due to the predominantly blue illumination obtained from a clear sky.

These pixel values shall be analysed and parameters estimated for both achromatic and colour reflection models.

7.4.1 Achromatic Analysis

This section uses simulated annealing to estimate the parameters of the achromatic reflection model for each of the material samples. For this purpose the image sequence has been converted to a greyscale sequence. The cooling schedule used for the analysis of this data is that used in Section 7.3.3. The resulting estimated parameter values are shown in Table 7.14.

Table 7.14: *Estimated achromatic model parameters for material samples.*

Sample number	Estimated Parameters			χ^2 of result	χ^2 evaluations
	R	ρ_s	σ_α		
1	0.44	0.28	0.00	6179	618
2	0.51	0.51	0.09	4535	629
3	0.66	0.51	0.01	6253	620
4	0.17	0.50	0.07	6813	570
5	0.61	0.83	0.13	5647	608
6	0.30	0.79	0.18	49817	479
7	0.36	0.50	0.00	1414	727
8	0.17	0.50	0.01	510	722

The correct parameter values cannot be determined without thorough analysis of the reflection characteristics of each material sample. However, the results are broadly consistent with expectations. In the case of the painted samples (3, 7 and 8) the estimated parameters differ only in the total reflectance, which would be expected since only the colour of the paint differs. The mild steel samples (4, 5 and 6) do exhibit increasing values of roughness, though the polished steel sample would be reasonably expected to have a higher value of ρ_s than that estimated. In general, the metallic samples have estimated ρ_s values higher than those of the other, non-metallic materials.

7.4.2 Colour Analysis

This section uses simulated annealing to estimate the parameters of the colour reflection model for each of the eight material samples. Here the cooling schedule used for the analysis is that used in

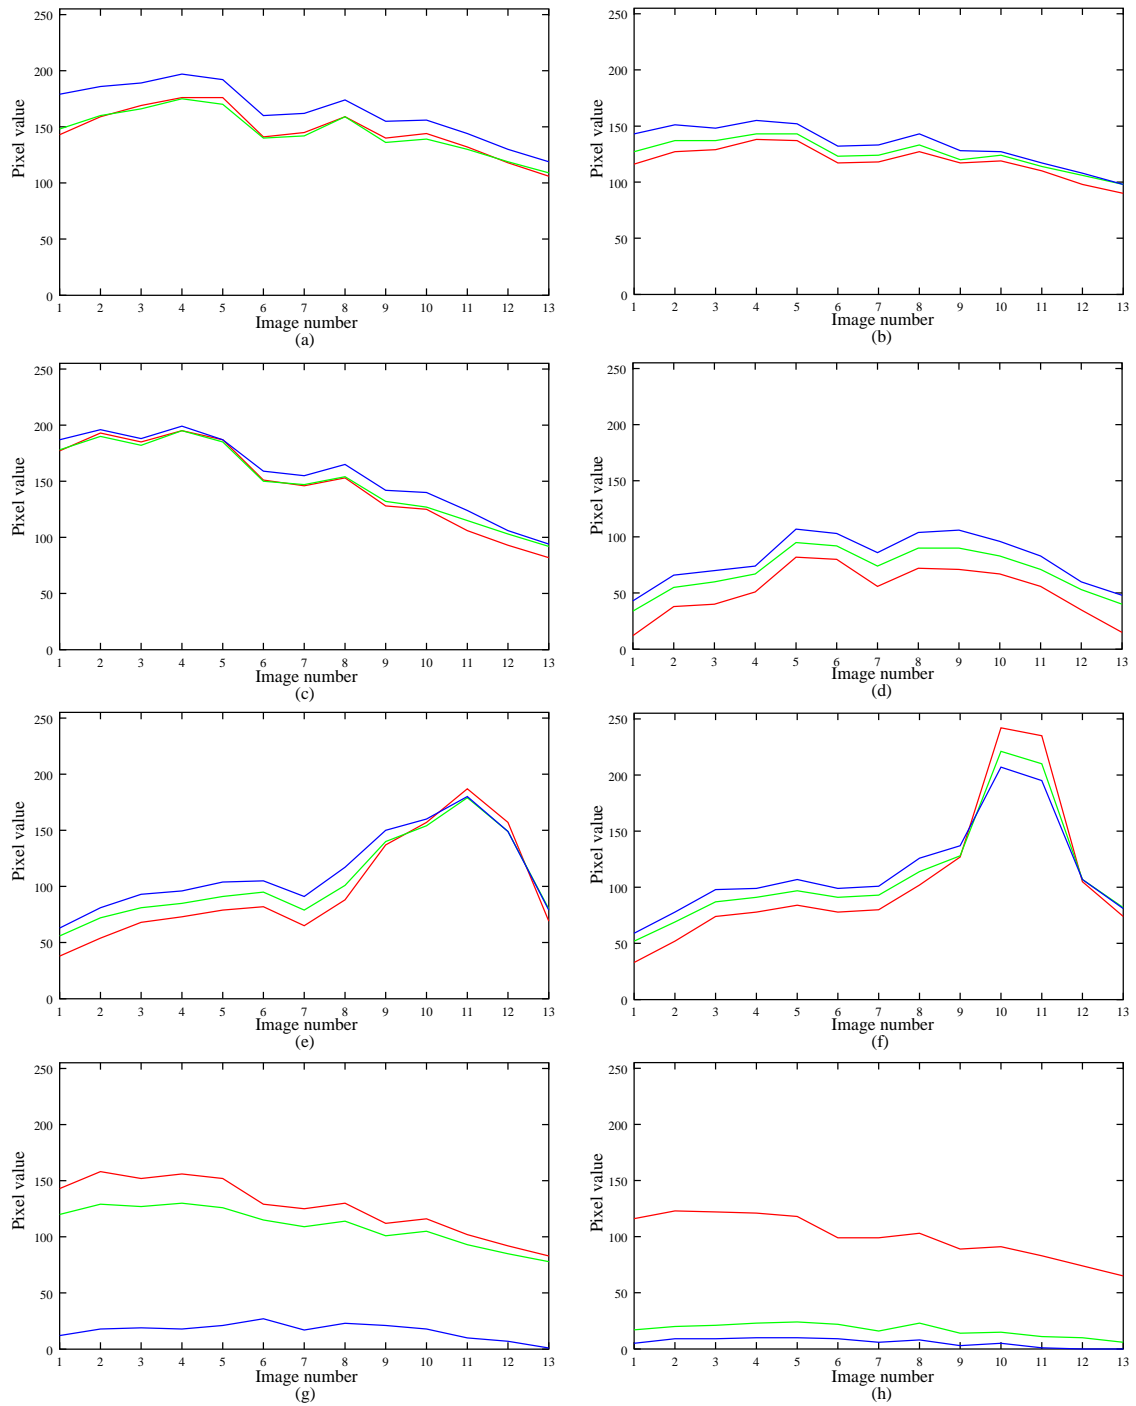


Figure 7.12: Values of a single pixel within the projection of each material sample across the image sequence: (a) sample 1; (b) sample 2; (c) sample 3; (d) sample 4; (e) sample 5; (f) sample 6; (g) sample 7; (h) sample 8.

Section 7.3.4. The resulting estimated parameter values are shown in Table 7.15.

Table 7.15: *Estimated colour model parameters for material samples.*

Sample number	Estimated Parameters					χ^2 of result	χ^2 evaluations
	C_{red}	C_{green}	C_{blue}	ρ_s	σ_α		
1	0.46	0.45	0.46	0.30	0.05	7974	1403
2	0.55	0.58	0.54	0.52	0.00	4472	1433
3	0.61	0.61	0.56	0.46	0.01	6536	1383
4	0.21	0.24	0.25	0.47	0.06	6692	1425
5	0.63	0.60	0.60	0.80	0.13	8111	1428
6	0.37	0.36	0.34	0.78	0.17	45225	1221
7	0.56	0.48	0.06	0.50	0.03	3155	1498
8	0.44	0.08	0.02	0.50	0.04	1325	1575

The predicted pixel values for each sample across the image sequence are shown in Figure 7.14. These have been calculated using the estimated model parameters found by simulated annealing. It can be seen that there is, in most cases, a strong similarity between the actual pixel values shown in Figure 7.12. Only in samples 3 and 4 has the model had difficulty in matching the actual pixel values. The estimated ρ_s and σ_α values for this colour reflection model exhibit a similarity to those found for the achromatic model. This shows a consistency in the results obtained by this method. It can also be seen that the method has correctly estimated the colour of the painted samples, as well as the dominance of the blue pixel values in each case. However, in the case of the pale blue paper (sample 2) the system has been unable to discern a colour difference. As stated above, the correct parameter values for each sample can only be determined by thorough analysis of the materials. However, Figure 7.14 shows that the estimated parameters are sufficient for an accurate reproduction of the measured data. A reconstruction of the image sequence, created using the estimated parameter values, is shown in Figure 7.13.

These results show some promise for the proposed method of reflection model parameter estimation. Failure to accurately model the observed data is due to inadequacies in the models. Each of the models employed by the method presents an approximation of the process involved. As such, the accuracy of the parameter estimation is limited by the accuracy of the models used for camera response, illumination and surface reflection.

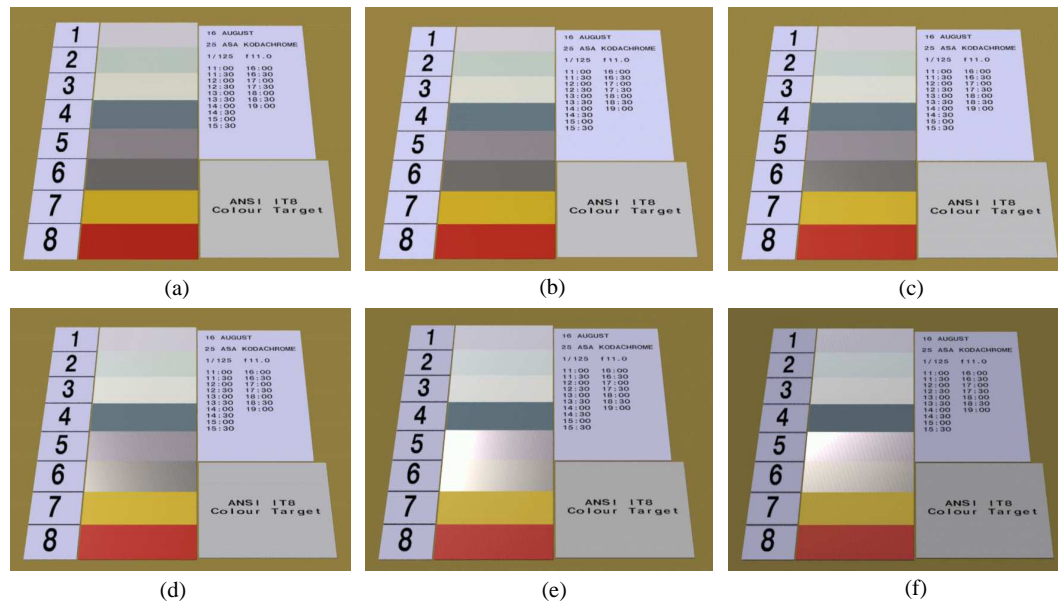


Figure 7.13: *Reconstruction of the real image sequence using estimated parameters: (a) 12:00; (b) 13:00; (c) 14:00; (d) 15:00; (e) 16:00; (f) 17:00.*

7.5 Summary

This chapter has examined the performance of two regression techniques on synthetic and real image sequences. For synthetic image sequences, which provide the most suitable conditions for parameter estimation, the Levenberg-Marquardt regression method has been shown to have serious deficiencies. These have been due to both the characteristics of the method and of the application in which it has been used. The failings of this method in the case of simple models with two and three parameters have been clearly demonstrated.

It has been shown that the method of simulated annealing provides consistent results, provided that a suitable cooling schedule is devised. The effectiveness of this method has been demonstrated in the case of simulated image sequences and reflection models with two, three and five parameters. The method has been shown to be successful in those cases where traditional regression methods have failed.

In the case of a real image sequence, simulated annealing has been used to estimate the parameters of both an achromatic and a colour reflection model. There has been some consistency in the results for each of these reflection models, though the accuracy of the values cannot be

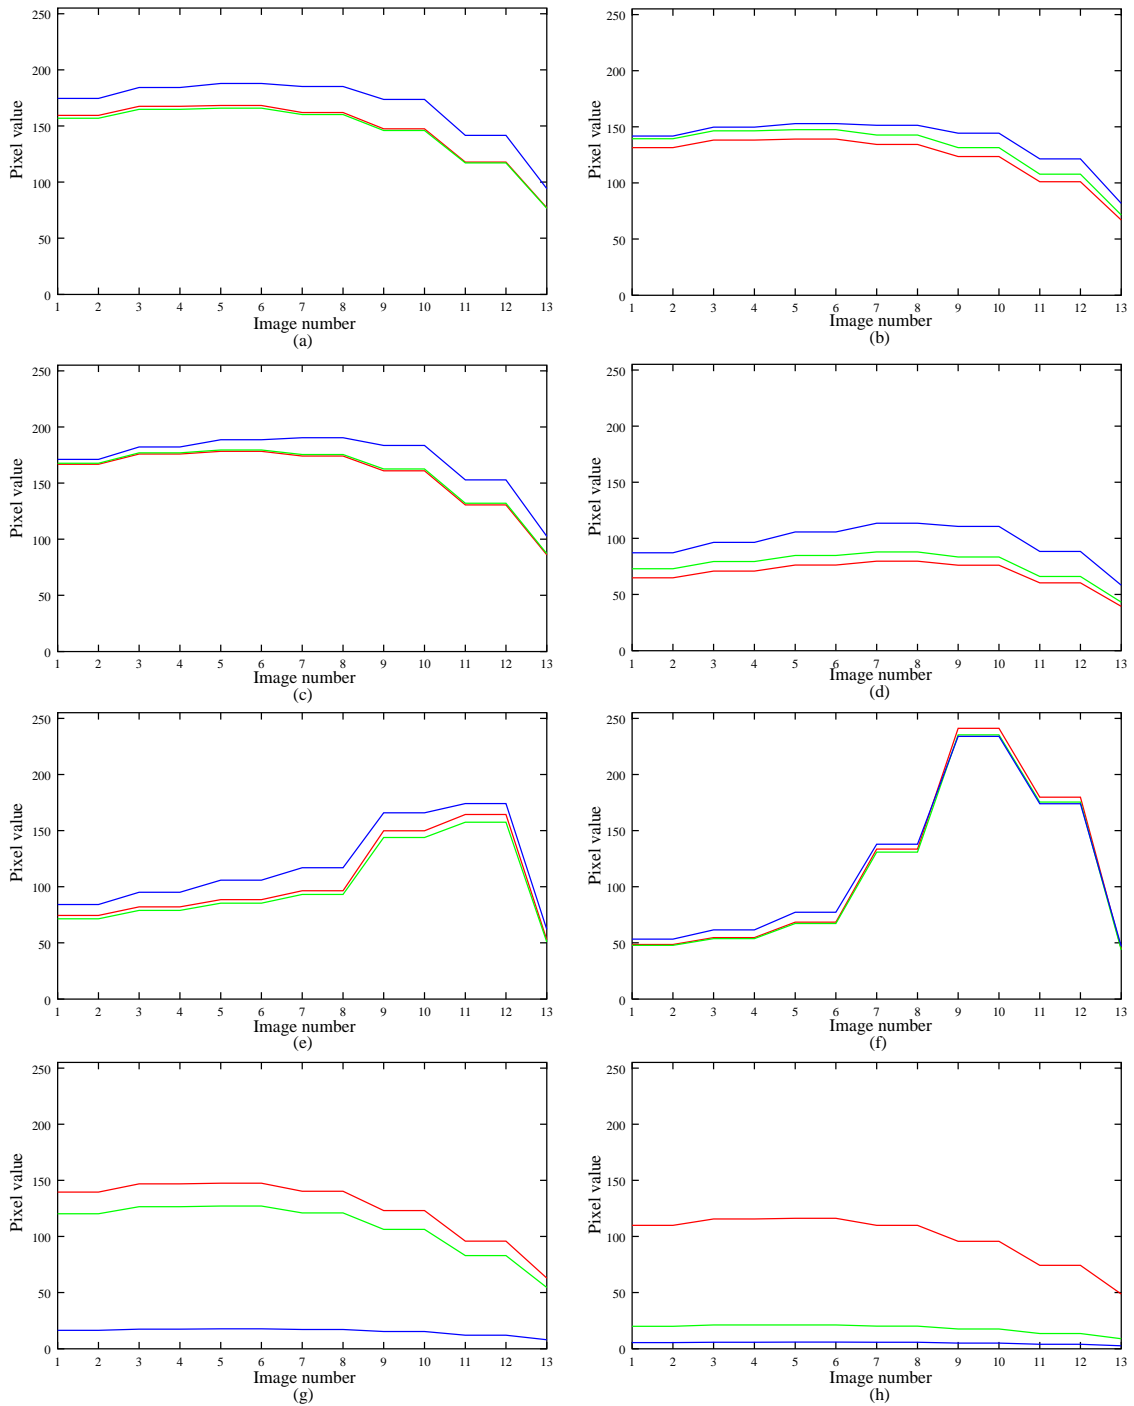


Figure 7.14: Pixel values for each material sample modelled using the estimated colour reflection model parameters: (a) sample 1; (b) sample 2; (c) sample 3; (d) sample 4; (e) sample 5; (f) sample 6; (g) sample 7; (h) sample 8.

evaluated without using material samples of known reflectance. However, the results show that the estimated colour reflection model parameters are suitable for an accurate simulation of the observed data. Such results would enable the simulation of the observed scene in differing illumination conditions and from alternative camera locations. This type of simulation could be used as a virtual environment depicting the observed scene.

Chapter 8

Conclusions

The aim of this thesis has been to develop a method for estimating a model of reflection for an object observed under natural illumination by a static camera. Further to this, it has been required that the results obtained are quantitative, such that they may be used for subsequent lighting calculations or simulations. Such results may be used for a variety of applications. These include object recognition, material analysis, or automated virtual environment generation from existing natural scenes. This latter application has been the subject of some study. It has been noted that the apparent realism of a virtual environment is not simply achieved by accurate modelling of object shape [6, 26]. Realism is enhanced through the use of accurate surface reflection models which provide a sense of texture, gloss and colour.

Chapter 3 has shown how the reflection characteristics of a surface may be encapsulated in a functional model. Such computational models allow the approximation of a bidirectional reflectance-distribution function (BRDF) for a surface. It has been shown that many such models exist, each formulated for a particular application or surface type. One model has been selected for evaluation in this thesis, that being the model due to Ward [141]. It has been shown that this model is applicable to many material types and that it has the advantage of being computationally convenient. Three variations of this basic reflection model have been considered, each with an increasing number of parameters to be estimated, and accounting for both achromatic and colour reflection.

This thesis has considered natural illumination consisting of two main sources: direct il-

lumination from the sun and diffuse illumination from the sky. Many computer graphics and machine vision applications have neglected the complexity of these sources. Typically the sun is modelled inaccurately as a point source. Illumination from the sky is often either neglected, substituted by an ambient term, or assumed to be a hemispherical source of uniform brightness. Neither of these naive simplifications adequately describes the complexity of natural illumination. Chapter 4 has presented physically-based models of sun and sky light which have been developed by the solar energy research community. As such, they accurately model the distribution and magnitude of radiance from both the sun and sky. It has also been shown that, by consideration of the scattering processes which light undergoes as it passes through the atmosphere, the spectral distribution of light may be modelled. Models have been presented which describe how natural illumination changes with time, terrestrial location, and weather conditions. Methods have also been compared which allow these illumination models to be used in the context of machine vision.

This research has used images obtained by a camera to measure surface reflection. In order that measurements can be made from the values of image pixels, it has been necessary to derive a sensor model suitable for this application. Such a model has been presented in Chapter 5. This has considered the camera optics required to project an image and the process by which an image is discretised to form a colour digital image. The use of a radiometric camera model allows image pixel values to be interpreted in the context of a surface reflecting incident irradiance. Measurements made from image data can, therefore, be used in the formulation of a BRDF model for an observed surface.

Chapter 6 has shown that, by considering a sequence of images, taken over a period of time, a scene may be observed in a variety of illumination conditions. This is due to the changing nature of daylight illumination over time. It has been shown that the temporal change in image pixel values, over the duration of an image sequence, provides sufficient information for the surface reflection characteristics to be modelled. It has been shown that such analysis may be performed using a single pixel, or multiple pixels, within each image of a sequence.

The estimation of reflection model parameters from observed image data has been achieved through the use of regression techniques. Such methods attempt to find model parameters which reduce the least-squares difference between observed data and a model. This difference is evaluated using a χ^2 metric. It has been shown that minimisation of this χ^2 function will yield the optimal reflection model parameters, in the least-squares sense. The use of the χ^2 metric also

provides a tolerance of image noise due to sensor error.

Two methods of regression have been considered. The first has been the Levenberg-Marquardt method which is commonly used for parameter estimation. The second method is the probabilistic approach of simulated annealing. The use of the latter method of regression is novel in this area of study. These two methods have been selected by considering the nature of the reflection models. The expected performance of each method has been analysed by considering a large number of synthetic image sequences. These have allowed the performance of each technique to be assessed in precisely controlled conditions.

The presented methods have required *a priori* knowledge of scene geometry, camera parameters and illumination conditions for each image in a sequence. This has been required so that the free parameters of the complete model, used to predict image pixel values, are those of the selected reflection model alone. Whilst these are significant assumptions, it has been shown that parameters other than those of the reflection model may be estimated using existing methods. In particular, scene geometry may be determined using a number of shape recovery techniques.

8.1 Performance

Chapter 7 has presented results for the three variants of the Ward reflection model that have been considered. Each of these models has had an increasing number of parameters to be estimated. The performance of the two proposed methods of regression has been evaluated using synthetic image sequences. These have allowed a thorough comparison of the techniques to be made for a variety of sequences captured under precisely known conditions.

It has been shown that the Levenberg-Marquardt method is not suitable for this application. Chapter 6 showed that the χ^2 function for some examples contained local minima corresponding to non-optimal parameters. Such parameters do not yield the best, least-squares, fit between measured and modelled data. The gradient descent approach will often identify these local minima as solutions. It has also been shown that, even in the simple cases where only a global optimal minimum exists, the accuracy of the estimated parameters is dependent upon the initial parameter estimates. Therefore, this technique requires an estimate of the reflection model parameters to be made that is sufficiently close to the optimal values for accurate results to be obtained. Such initial

estimates are not available without prior analysis of the image data.

The method of simulated annealing has shown itself to be useful for this application. Whilst it requires a significantly increased number of computations to be made, the method shows none of the failings of the Levenberg-Marquardt method. The technique is not misled by local, non-optimal, minima and provides consistent results independently of the initial starting parameters. Despite the increased number of χ^2 function evaluations that are required, it has been shown that the computational complexity of the method is less than that of a naive brute force search of the parameter space. It has been shown that, given a suitable cooling schedule, simulated annealing is able to consistently estimate parameters for the achromatic and colour reflection models. This has been shown through the analysis of synthetic image sequences.

The analysis of real images presents a significant test of any machine vision system. In this case, a single clear sky sequence showing a variety of eight planar surface samples has been used. This sequence has been analysed and model parameters estimated for both the achromatic and colour reflection models. The estimated parameters for each reflection model have been shown to be consistent for each of the eight surface samples. These estimated parameter values have correctly indicated surface roughness, surface gloss and material colour. However, the accuracy of the estimated parameters cannot be determined without the use of calibrated materials.

8.2 Future Work

The current method assumes *a priori* knowledge of weather conditions, camera parameters and surface orientation. These assumptions represent a significant amount of information which is required for each image of a sequence. However, each of these factors are merely parameters of the various models used to describe illumination, camera response and surface orientation. Given that the method of simulated annealing scales well with an increasing number of parameters, it might be possible to estimate these in the regression process. For example, the method could be extended to a situation in which surface normal is unknown. In this case, the regression would estimate not only the parameters of the reflection model but also the orientation of the surface.

Chapter 3 has shown that there is no single reflection model which is suitable for characterising the reflection from all types of surface material. This is especially true for the many natural

surfaces, such as vegetation or sand, which may be experienced in a terrestrial scene. Given that the regression methods developed here are independent of the selected reflection models, it would be possible to select the model most suited to the observed scene. In this case the system would select the model, as well as the associated parameters, which yield a minimum least-squares difference. However, there could be many problems associated with such a generalised approach. Primarily, the effect on the regression method of a transition between reflection models, possibly resulting in a change of parameter space, would need to be understood. The use of many reflection models could result in a global minimum which is not the desired solution.

The presented methods have considered static scenes observed by a static camera. However, a moving camera or animated scene would be expected to provide an increased variation in image pixel values. As has been demonstrated, an observed highlight serves to distinguish between matte and glossy surfaces. Changing the orientation of a surface with respect to the camera would provide an increased probability of observing such a highlight. This improvement would not require any change to the regression method or to the dimension of the parameter space, provided that camera and object position are recorded for each frame.

The observed shading of a surface is not due only to direct illumination from light sources but also interreflection from other surfaces. All the scenes considered in this work have been such that interreflection could not take place. It may be possible to extend the methods to account for interreflection using global illumination techniques. However, this presents a considerably more complex problem than that presented here. In the case of an occluded surface reflecting light onto a visible surface, it would not be possible to determine the quantity of light being mutually reflected, if the reflectivity of both surfaces is unknown.

The method of simulated annealing requires a significant amount of computing time in order to converge to a solution. This is due to the repeated evaluation of the χ^2 function for each hypothesised set of parameter values. Each χ^2 evaluation requires the irradiance upon the sensor to be evaluated for each considered pixel in the sequence. This is clearly a computationally intensive task, the performance of which could be improved with further investigation. One approach would be to parallelise the algorithm. In the simulated annealing case, each vertex of the simplex could be evaluated independently. Similarly, the calculation of the difference between predicted and measured pixel values, for each image in the sequence, could be distributed.

8.3 Closing Comments

This thesis has shown that the estimation of reflection model parameters is possible for surfaces visible in images of naturally illuminated scenes. This is a significant departure from the study of laboratory scenes, in which the problem of modelling reflectance is usually much simpler. This work has shown that existing, physical, models of natural illumination may be successfully employed in machine vision. These models allow quantitative results to be obtained, which can then be used for a variety of applications.

Appendix A

Radiometry and Photometry for Machine Vision

To simplify the sharing of data and methods between scientific communities a standard set of metrics and definitions are required. Where possible this thesis has used recognised notation as defined by the Illumination Engineering Society [54]. This appendix defines some of the terms presented in this thesis.

Radiometry is the science of the physical measurement of electromagnetic radiation. A radiometric measurement of *radiant energy* is expressed in Joules [J]. The amount of energy at each wavelength of light can be measured using a spectroradiometer. The resulting measurements provide the spectrum of the light source. Photometry, however, is the psychophysical measurement of the visual sensation experienced by a human observer to the light. Pierre Bouguer (1760), founder of the field of photometry, first noted that a human observer could not provide a quantitative description of perceived brightness but was able to compare the brightness of two light sources. Since the human visual system has varying response to wavelength, sources of equal radiant energy but differing spectrum may have different perceived brightnesses. This relative sensitivity to wavelength can be plotted on a curve termed the *spectral luminous efficacy curve*, see Figure A.1. Brightness, or *luminous energy*, is measured in units of talbots.

The rate of flow of radiant energy, *radiant power*, between two points is measured in Joules per second or Watts [W]. The photometric equivalent is the rate of flow of luminous energy, *lu-*




Figure A.1: *Spectral luminous efficacy curve.*

minous power, is measured in talbots per second or lumens. Since radiometry is more physically fundamental than photometry, radiometric quantities may be computed from their photometric equivalents. This is possible by considering the *efficacy* of the light at each wavelength. Efficacy is measured in lumens per Watt and defines the energy per lumen for light of a particular wavelength. For example, considering the visible wavelengths of light, the radiant power, Φ , of a source with luminous power L and luminous efficacy E is;

$$\Phi = \int_{\lambda=380\text{nm}}^{770} \frac{L(\lambda)}{E(\lambda)} d\lambda. \quad (\text{A.1})$$

In measuring radiant power, or flux, it is necessary to consider the direction of flow. This is given as solid angle of directions measured in steradians [sr]. A solid angle is the three dimensional equivalent of angle in plane geometry. A solid angle, ω , is measured in terms of the area on a sphere intercepted by a cone whose apex is at the sphere's centre. A unit solid angle intercepts an area equal to the square of the sphere's radius, r [24]. A hemisphere, therefore, has a solid angle of $4\pi r^2 / 2r^2 = 2\pi$ sr. A differential solid angle of directions, $d\omega$, may be termed a *ray*.

The projected area of a surface is the apparent area of the surface seen by an observer from a particular direction. This projected area, dA_{-} , is the surface area dA multiplied by the cosine of the angle θ which the surface normal makes with the observer, $dA_{-} = dA \cos \theta$ [103]. The projected area of a surface is, therefore, dependent upon the relative orientation with the observer. The solid angle subtended by a surface about a particular point may be derived in terms of the

projected area. A surface of projected area dA_{\perp} , at a distance r from a point, subtends a solid angle $d\omega$:

$$d\omega = \frac{dA_{\perp}}{r^2} = \frac{dA \cos \theta}{r^2}. \quad (\text{A.2})$$

The following are radiometric quantities which may be derived from the basic units presented above. Each is accompanied, where applicable, by the corresponding photometric term.

Radiant intensity [W sr^{-1}] : The power per unit solid angle radiated about a particular direction.

The corresponding photometric quantity is *luminous intensity* [Candela].

Radiance [$\text{W m}^{-2} \text{sr}^{-1}$] : Radiance, L , is the power per unit projected area perpendicular to the ray per unit solid angle in the direction of the ray. Radiance has two useful properties [12]:

1. The radiance in the direction of a ray remains constant as it propagates along the ray. This is valid provided there are no losses due to scattering or absorption.
2. The response of a sensor is proportional to the radiance of the surface visible to the sensor.

The corresponding photometric quantity is *luminance* [Nit].

Irradiance [W m^{-2}] : The total radiant energy per unit area incident on a surface of fixed orientation from the hemisphere of incident directions is termed irradiance, E . Irradiance is calculated by integrating incident radiance, L_i , over the visible hemisphere, Ω ;

$$E = \int_{\Omega} L_i \cos \theta_i d\omega, \quad (\text{A.3})$$

where θ_i is the angle of incidence relative to the zenith. The corresponding photometric quantity is *illuminance* [Lux].

Radiosity [W m^{-2}] : Whereas irradiance is due to incident radiance, radiosity, B , is the energy per unit area that leaves a surface. The corresponding photometric quantity is *luminosity* [Lux].

Bibliography

- [1] R. Bajcsy, S. W. Lee, and A. Leonardis. Color image segmentation with detection of highlights and local illumination induced by interreflections. In *Proceedings of International Conference on Pattern Recognition*, pages 785–790. IEEE Computer Society Press, June 1990.
- [2] R. Baribeau, M. Rioux, and G. Godin. Color reflectance modeling using a polychromatic laser range finder. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):263–269, 1992.
- [3] D. M. Bates and D. G. Watts. *Nonlinear regression and its applications*. John Wiley & Sons, 1988.
- [4] P. Beckmann and A. Spizzichino. *The Scattering of Electromagnetic Waves from Rough Surfaces*. Pergamon Press, 1963.
- [5] P. R. Bevington. *Data reduction and error analysis for the physical sciences*. McGraw-Hill, New York, second edition, 1994.
- [6] G. Bishop and H. Fuchs. Research directions in virtual environments: Report of an NSF invitational workshop. *Computer Graphics*, 26(3):153–177, 1992.
- [7] J. F. Blinn. Models of light reflection for computer synthesized pictures. *Computer Graphics (SIGGRAPH'77)*, 11(2):192–198, July 1977.
- [8] I. O. Bohachevsky, M. E. Johnson, and M. L. Stein. Generalized simulated annealing for function optimization. *Technometrics*, 28(3):209–217, 1986.
- [9] D. H. Brainard and B. A. Wandell. Analysis of the retinex theory of color vision. *Journal of the Optical Society of America A*, 3(10):1651–1661, 1986.

- [10] M. H. Brill and G. West. Chromatic adaption and color constancy: A possible dichotomy. *COLOR Research and Applications*, 11:196–204, 1986.
- [11] B. J. Buratti. Photometry of rough planetary surfaces: The role of multiple scattering. *Icarus*, 64:320–328, 1985.
- [12] M. F. Cohen and Wallace J. R. *Radiosity and Realistic Image Synthesis*. Academic Press, 1993.
- [13] E. N. Coleman and R. Jain. Obtaining 3D shape of textured and specular surfaces using four-source photometry. *Computer Vision, Graphics, and Image Processing*, 18(4):309–328, 1982.
- [14] R. L. Cook and K. E. Torrance. A reflection model for computer graphics. *Computer Graphics*, 15(3), 1981.
- [15] A. Corana, M. Marchesi, C. Martini, and S. Ridella. Minimizing multimodal functions of continuous variables with the simulated annealing algorithm. *ACM Transactions of Mathematical Software*, 13(3):262–280, September 1987.
- [16] K. J. Dana, S. K. Nayer, B. van Ginneken, and J. J. Koenderink. Reflectance and texture of real-world surfaces. Technical Report CUCS-048-96, Columbia University, Department of Computer Science, Columbia University, New York, 1996.
- [17] Commision Internationale de l'Éclairage. Recommendations for the integrated irradiance and the spectral distribution of simulated solar radiation for testing purposes, 1972. CIE Publication 20.
- [18] Commision Internationale de l'Éclairage. Spatial distribution of daylight - luminance distribution of various reference skies, 1994. CIE Publication 110.
- [19] Y. Dobashi, K. Kaneda, H. Yamashita, and T. Nishita. Method for calculation of sky light luminance aiming at an interactive architectural design. In *Proceedings of EUROGRAPHICS'96*, pages 109–118, 1996.
- [20] N. Dodd. Multispectral texture synthesis using fractal concepts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9:703–707, 1987.

- [21] M. D'Zmura and P. Lennie. Mechanisms of color constancy. *Journal of the Optical Society of America A*, 3(10):1662–1672, 1986.
- [22] W. Egan and T. Hilgeman. Retroreflectance measurements of photometric standards and coatings. *Applied Optics*, 15:1845–1849, 1976.
- [23] G. D. Finlayson, B. V. Funt, and K. Barnard. Color constancy under varying illumination. In *Proceedings of International Conference on Computer Vision*, pages 720–725. IEEE Computer Society Press, June 1995.
- [24] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Huges. *Computer Graphics: Principles and Practice*. Addison-Wesley, second edition, 1996.
- [25] B. V. Funt, M. S. Drew, and J. Ho. Color constancy from mutual reflection. *International Journal of Computer Vision*, 6(1):5–24, 1991.
- [26] A. Gagalowicz. Tools for advanced telepresence systems. *Computers and Graphics*, 19(5):73–88, 1995.
- [27] R. Gershon. *The Use of Color in Computational Vision*. PhD thesis, Department of Computer Science, University of Toronto, 1987.
- [28] D. Gibbens, M. J. Brooks, and W. Chojnacki. Light source direction from a single image: a performance analysis. *The Australian Computer Journal*, 23(4):165–174, 1991.
- [29] G. Gillette and T. Kusuda. A daylighting computation procedure for use in DOE-2 and other dynamic building energy analysis programs. *Journal of the Illumination Engineering Society*, 12(2):78–85, January 1983.
- [30] A. S. Glassner. *An Introduction to Ray Tracing*. Academic Press, 1989.
- [31] A. S. Glassner. *Principles of digital image synthesis*, volume 2. Morgan Kaufmann, 1995.
- [32] A. S. Glassner. *Principles of digital image synthesis*, volume 1. Morgan Kaufmann, 1995.
- [33] N. S. Goel and R. L. Thompson. Inversion of vegetation canopy reflectance models for estimating agronomic variables. *Remote Sensing of the Environment*, 14:77–111, 1984.
- [34] W. E. L. Grimson. Binocular shading and visual surface reconstruction. *Computer Vision, Graphics, and Image Processing*, 28(1):19–43, 1984.

- [35] A. Grynberg. Validation of Radiance. Technical Report LBID 1575, Lawrence Berkley Laboratory, July 1989.
- [36] R. Hall. *Illumination and Color in Computer Generated Imagery*. Springer-Verlag, 1988.
- [37] B. Hapke. Bidirectional reflection spectroscopy: 3. Correction for macroscopic roughness. *Icarus*, 59:41–59, 1984.
- [38] B. Hapke. Coherent backscatter and the radar characteristics of outer planet satellites. *Icarus*, 88:407–417, 1990.
- [39] T. Hashimoto, H. Kato, K. Sato, S. Inokuchi, and K. Moriwaki. Recognition of material types and interreflection using color images. *Systems and Computers in Japan*, 23(12):727–735, 1992.
- [40] D. I. Havelock. Geometric precision in noise free digital images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(10):1065–1075, October 1989.
- [41] X. D. He, P. O. Heynen, R. L. Phillips, K. E. Torrance, D. H. Salesin, and D. P. Greenberg. A fast and accurate light reflection model. *Computer Graphics (SIGGRAPH'92)*, 26(2):253–254, July 1992.
- [42] X. D. He, K. E. Torrance, F. X. Sillion, and D. P. Greenberg. A comprehensive physical model for light reflection. *Computer Graphics (SIGGRAPH'91)*, 25(4):175–186, July 1991.
- [43] G. Healey and T. O. Binford. Local shape from specularity. *Computer Vision, Graphics, and Image Processing*, 42:62–86, 1988.
- [44] G. Healey and R. Kondepudy. Modeling and calibrating CCD cameras for illumination insensitive machine vision. In *Proceedings of Optics, Illumination, and Image Sensing for Machine Vision VI*. Society of Photo Optical Instrumentation Engineers, 1991.
- [45] G. Healey and R. Kondepudy. Radiometric CCD camera calibration and noise estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(3):267–276, March 1994.
- [46] S. T. Henderson and D. Hodgkiss. The spectral energy distribution of daylight. *British Journal of Applied Physics*, 15(947), 1964.

- [47] J. Ho, B. V. Funt, and M. S. Drew. Separating a color signal into illumination and surface spectral components. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(10):966–977, 1990.
- [48] R. G. Hopkinson, P. Petherbridge, and J. Longmore. *Daylighting*. Heinemann, 1966.
- [49] B. J. P. Horn. *Shape from shading: A method for obtaining the shape of a smooth opaque object from one view*. PhD thesis, Department of Electrical Engineering, MIT, Cambridge, MA, 1970.
- [50] B. K. P. Horn and M. J. Brooks. *Shape from shading*. MIT Press, Cambridge, MA, 1989.
- [51] B. K. P. Horn and R. W. Sjoberg. Calculating the reflectance map. *Applied Optics*, 18(11):1170–1779, 1979.
- [52] K. Ikeuchi. Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3(6):661–669, 1981.
- [53] K. Ikeuchi and K. Sato. Determining reflectance properties of an object using range and brightness images. *IEEE Transactions of Pattern Analysis and Machine Intelligence*, 13(11):1139–1153, 1990.
- [54] Illumination Engineering Society. Nomenclature and definitions for illumination engineering. Technical Report ANSI/IES RP-16-1986, American National Standards Institute, 1986.
- [55] M. Inakage. An illumination model for atmospheric environments. In *New Advances in Computer Graphics (Proceedings of Computer Graphics International 1989)*, pages 533–548. Springer-Verlag, 1989.
- [56] D. B. Judd, D. L. MacAdam, and G. Wyszecki. Spectral distribution of typical daylight as a function of correlated color temperature. *Journal of the Optical Society of America*, 54(1031), 1964.
- [57] B. Kamgar-Parsi and B. Kamgar-Parsi. Evaluation of quantization error in computer vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(9):929–940, September 1989.

- [58] K. F. Karner, H. Mayer, and M. Gervautz. An image based measurement system for anisotropic reflection. In *Proceedings of EUROGRAPHICS'96*, 1996.
- [59] G. Kay and T. Caelli. Inverting an illumination model from range and intensity maps. *CVGIP Image Understanding*, 59(2):183–201, 1994.
- [60] G. Kay and T. Caelli. Estimating the parameters of an illumination model using photometric stereo. *CVGIP Graphical Models and Image Processing*, 57(5):365–388, 1995.
- [61] D. S. Kimes. Modeling the directional reflectance from complete homogeneous vegetation canopies with various leaf orientation distributions. *Journal of the Optical Society of America*, 1:725–737, 1984.
- [62] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. Technical Report RC 9355, IBM Research Report, 1982.
- [63] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 1983.
- [64] R. Kittler. Standardisation of outdoor conditions for the calculation of daylight factor with clear skies. In *Proceedings of CIE Conference on Sunlight*, Newcastle, 1965.
- [65] R. V. Klassen. Modeling the effect of the atmosphere on light. *ACM Transactions on Graphics*, 6(3):215–237, July 1987.
- [66] G. J. Klinker. *A physical approach to color image understanding*. A K Peters, 1993.
- [67] J. J. Koenderink and A. J. van Doorn. Geometrical modes as a general method to treat diffuse interreflections in radiometry. *Journal of the Optical Society of America*, 73(6):843–850, 1983.
- [68] J. J. Koenderink, A. J. van Doorn, and M. Stavridi. Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *Proceedings of European Conference on Computer Vision*, pages 28–39, 1996.
- [69] C. Kolb, D. Mitchell, and P. Hanrahan. A realistic camera model for computer graphics. *Computer Graphics (SIGGRAPH'95)*, 29(8):317–324, 1995.
- [70] J. Krochmann and M. Seidl. Quantitative data on daylight illumination engineering. *Lighting Research and Technology*, 6(3), 1974.

- [71] Lambert. *Photometria sive de mensura et gradibus luminis, colorum et umbrae*, 1760.
- [72] E. H. Land and J. J. McCaen. Lightness and the retinex theory. *Journal of the Optical Society of America*, 61:1–11, 1971.
- [73] M. S. Langer and S. W. Zucker. Shape-from-shading on a cloudy day. *Journal of the Optical Society of America A*, 11(2):467–478, 1994.
- [74] H.-C. Lee, E. J. Breneman, and C. P. Shulte. Modeling light reflection for color computer vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(4):402–409, 1990.
- [75] K. M. Lee and C.-C. J. Kuo. Shape from shading with a generalized reflectance map model. *CVGIP Computer Vision and Image Understanding*, 67(2):143–160, 1997.
- [76] S. W. Lee. *Understanding of Surface Reflection in Computer Vision by Color and Multiple Views*. PhD thesis, GRASP Laboratory, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, 1991.
- [77] P. J. Littlefair. A comparison of sky luminance models with measured data from Garston, United Kingdom. *Solar Energy*, 53(4):315–322, 1994.
- [78] R. C. Love and N. D. Efford. Surface reflectance model estimation from daylight illuminated image sequences. In *Proceedings of the British Machine Vision Conference*, pages 317–326. BMVA, 1995.
- [79] J. Lu and J. Little. Reflectance function estimation and shape recovery from image sequence of a rotating object. In *Proceedings of IEEE International Conference on Computer Vision*, pages 80–86. IEEE Computer Society Press, June 1995.
- [80] L. T. Maloney and B. A. Wandell. Color constancy: A method for recovering surface spectral reflectance. *Journal of the Optical Society of America A*, 3(1):29–33, 1986.
- [81] J. Mardaljevic. Validation of a lighting simulation program under real sky conditions. *International Journal of Lighting Research and Technology*, 27(4):181–188, 1995.
- [82] A. G. Maristany, P. K. Lebow, C. C. Bruner, D. A. Butler, and K. W. Funck. Classifying wood-surface features using dichromatic reflection. In *Proceedings of SPIE Conference on Optics in Agriculture and Forestry*. SPIE, November 1992.

- [83] D. W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.
- [84] C. McCamy, H. Marcus, and J. Davidson. A color-rendition chart. *Journal of Applied Photographic Engineering*, 2(3), Summer 1976.
- [85] E. J. McCartney. *Optics of the atmosphere*. John Wiley & Sons, New York, 1976.
- [86] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.
- [87] W. Montgomery and R. Kohl. Opposition effect experimentation. *Optics Letters*, 5:546–548, 1980.
- [88] N. Mukawa. Optical-model-based analysis of consecutive images. *CVGIP Image Understanding*, 66(1):25–32, 1997.
- [89] J. F. Murray-Coleman and A. M. Smith. The automated measurement of BRDFs and their application to luminaire modeling. *Journal of the Illuminating Engineering Society*, Winter 1990.
- [90] V. Narasimhan and B. K. Saxena. Luminance and illumination from clear skies in the tropics. In *Proceedings of CIE Barcelona Session*, January 1971.
- [91] S. K. Nayer, X.-S. Fang, and T. Boult. Separation of reflection components using color and polarization. *International Journal of Computer Vision*, 21(3):163–186, 1997.
- [92] S. K. Nayer and K. Ikeuchi. Shape from interreflections. In *Proceedings of IEEE Conference on Computer Vision*, Osaka, December 1990. IEEE Computer Society Press.
- [93] S. K. Nayer, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Transactions on Robotics and Automation*, 6(4), 1990.
- [94] S. K. Nayer, K. Ikeuchi, and T. Kanade. Surface reflection: Physical and geometrical considerations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3(7):611–634, 1991.

- [95] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometrical considerations and nomenclature for reflectance. Technical Report BMS Monograph 160, National Bureau of Standards, October 1977.
- [96] T. Nishita and E. Nakamae. Continuous tone representation of three-dimensional objects illuminated by skylight. *Computer Graphics (SIGGRAPH'86)*, 20(4):125–132, August 1986.
- [97] C. Novak and S. A. Shafer. Anatomy of a color histogram. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 599–605. IEEE Computer Society Press, 1992.
- [98] C. L. Novak and S. A. Shafer. Supervised color constancy for machine vision. In *Proceedings of SPIE: Human Vision, Visual Processing and Digital Display II*, volume 1453, pages 353–368, March 1991.
- [99] C. L. Novak, S. A. Shafer, and R. G. Wilson. Obtaining accurate color images for machine vision research. In *Proceedings of SPIE: Perceiving, Measuring and Using Color*, volume 1250, pages 54–68. Society of Photo Optical Instrument Engineers, 1990.
- [100] Y. Ohta and Y. Hayashi. Recovery of illuminant and surface colors from images based on CIE daylight. In *Proceedings of European Conference on Computer Vision*, pages 269–280. Springer-Verlag, May 1994.
- [101] M. Oren and S. K. Nayer. Generalization of the Lambertian model and implications for machine vision. Technical Report CUCS-057-92, Department of Computer Science, Columbia University, New York, November 1992.
- [102] M. Oren and S. K. Nayer. Generalization of the lambertian model and implications for machine vision. *International Journal of Computer Vision*, 14:227–251, 1995.
- [103] S. N. Pattanaik. *Computational methods for global illumination and visualisation of complex 3D environments*. PhD thesis, Birla Institute of Technology and Science, Pilani, 1993.
- [104] A. P. Pentland. Finding the illuminant direction. *Journal of the Optical Society of America*, 72(4):448–455, 1982.
- [105] A. P. Pentland. Photometric motion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(9):879–890, 1991.

- [106] R. Perez, P. Ineichen, and R. Seals. Modelling daylight availability and irradiance components from direct and global irradiance. *Solar Energy*, 44(5):271–289, 1990.
- [107] R. Perez, R. Seals, and J. Michalsky. All-weather model for sky luminance distribution - preliminary configuration and validation. *Solar Energy*, 50(3):235–245, 1993.
- [108] B. T. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6), 1975.
- [109] B. Pinty, M. M. Verstraete, and R. E. Dickenson. A physical model for predicting bidirectional reflectances over bare soil. *Remote Sensing of Environment*, 27:273–288, 1989.
- [110] P. Poulin and A. Fournier. A model of anisotropic reflection. *Computer Graphics (SIGGRAPH'90)*, 24(4):273–281, August 1990.
- [111] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C - The Art of Scientific Computing*. Cambridge University Press, second edition, 1992.
- [112] D. Rapp. *Solar Energy*. Prentice-Hall, 1981.
- [113] D. Rapp and A. A. J Hoffman. On the relation between insolation and climatological variables-III. The relation between normal incidence solar intensity, total insolation, and weather at Fort Hood, Texas. *Energy Conversion*, 17(163), 1977.
- [114] J. C. Russ. Light scattering from fractal surfaces. *Journal of computer-assisted microscopy*, 5(2), 1993. 171-189.
- [115] V. V. Salomonson and W. E. Marlatt. Anisotropic solar reflectance over white sand, snow and stratus clouds. *Journal of Applied Meteorology*, 7:475–483, 1968.
- [116] Y. Sato and K. Ikeuchi. Temporal-color analysis of reflection. *Journal of the Optical Society of America A*, 11(11):2990–3002, 1994.
- [117] Y. Sato and K. Ikeuchi. Reflectance analysis under solar illumination. In *Proceedings of the Workshop on Physics-Based Modeling in Computer Vision*, pages 180–187. IEEE Computer Society Press, June 1995.
- [118] Y. Sato and K. Ikeuchi. Reflectance analysis for 3D computer graphics model generation. *CVGIP Graphical Models and Image Processing*, 58(5):437–451, 1996.

- [119] S. A. Shafer. Using color to separate reflection components. *COLOR research and applications*, 10(4):210–218, Winter 1985.
- [120] P. Shirley, C. Wang, and K. Zimmerman. Monte Carlo techniques for direct lighting calculations. *ACM Transactions on Graphics*, 15(1):1–36, January 1996.
- [121] R. Siegel and J. R. Howell. *Thermal Radiation Heat Transfer*. Hemisphere Publishing Corporation, third edition, 1992.
- [122] W. M. Silver. Determining shape and reflectance using multiple images. Master's thesis, Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA, 1980.
- [123] K. Solomon and K. Ikeuchi. Extracting the shape and roughness of specular lobe objects using 4 light photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(4):449–454, 1996.
- [124] E. M. Sparrow and R. D. Cess. *Radiation Heat Transfer*. Hemisphere Publishing Corporation, 1978.
- [125] K. Tadamura, E. Nakamae, K. Kaneda, M. Baba, H. Yamashita, and T. Nishita. Modelling of skylight and rendering of outdoor scenes. In *Proceedings of EUROGRAPHICS'93*, pages 189–200, 1993.
- [126] H. D. Tagare and R. J. P. deFigueiredo. Simultaneous estimation of shape and reflectance map from photometric stereo. *CVGIP Image Understanding*, 55(3), 1992. 275–286.
- [127] H. D. Tagare and R. J. P. deFigueiredo. A framework for the construction of reflectance maps for machine vision. *CVGIP Image Understanding*, 57(3):265–282, 1993.
- [128] A. H. Taylor and G. P. Kerr. The distribution of energy in the visible spectrum of daylight. *Journal of the Optical Society of America*, 31(947), 1941.
- [129] S. Tominaga. Surface identification using the dichromatic reflection model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13:658–670, 1991.
- [130] S. Tominaga. Dichromatic reflection models for a variety of materials. *COLOR research and applications*, 19(4):277–285, August 1994.
- [131] S. Tominaga. Multichannel vision system for estimating surface and illumination functions. *Journal of the Optical Society of America A*, 13(11):2163–2173, 1996.

- [132] S. Tominaga and B. A. Wandell. The standard surface reflectance model and illuminant estimation. *Journal of the Optical Society of America A*, 6(4):576–584, 1989.
- [133] S. Tominaga and B. A. Wandell. Component estimation of surface spectral reflectance. *Journal of the Optical Society of America A*, 7(2):312–317, 1990.
- [134] K. E. Torrance and E. M Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, 57(9):1105–1114, September 1967.
- [135] P. R. Tregenza. A simple mathematical model of illumination from a cloudy sky. *Lighting Resource and Technology*, 12(3):121–128, 1980.
- [136] P-S Tsai and M. Shah. Shape from shading using linear approximation. Technical Report CS-TR-92-24, Department of Computer Science, University of Central Florida, 1992.
- [137] R. Y. Tsai. Synopsis of recent progress on camera calibration for 3D machine vision. *Robotics Review*, 1989. MIT Press.
- [138] P. J. M. van Laarhoven and E. H. L. Aarts. *Simulated Annealing: Theory and Applications*. Kluwer Academic Publishers, 1987.
- [139] D. Vanderbilt and S. G. Louie. A Monte Carlo simulated annealing approach to optimization over continuous variables. *Journal of Computational Physics*, 56:259–271, 1984.
- [140] A. J. Verbiscer and J. Veverka. Scattering properties of natural snow and frost: Comparison with icy satellite photometry. *Icarus*, 88:418–428, 1990.
- [141] G. Ward. Measuring and modeling anisotropic reflection. *Computer Graphics (SIGGRAPH'92)*, 26(2):265–272, July 1992.
- [142] G. Ward. The RADIANCE lighting simulation and rendering system. *Computer Graphics (SIGGRAPH'94)*, pages 459–472, July 1994.
- [143] G. Ward, F. Rubinstein, and R. Clear. A ray tracing solution for diffuse interreflection. *Computer Graphics*, 22(4), August 1988.
- [144] A. H. Watt and M. Watt. *Advanced Animation and Rendering Techniques*. Addison-Wesley, 1992.

- [145] S. H. Westin, J. R. Avro, and K. E. Torrance. Predicting reflectance functions from complex surfaces. *Computer Graphics (SIGGRAPH'92)*, 26(2):255–264, July 1992.
- [146] N. Wittels, J. R. McClellan, K. Cushing, W. Howard, and A. Palmer. How to select cameras for machine vision. In *Proceedings of SPIE: Optics, Illumination and Image Sensing for Machine Vision III*, volume 1005, pages 44–53. Society of Photo Optical Instrument Engineers, 1988.
- [147] L. B. Wolff. Shape from polarization images. In *Proceedings of the IEEE Workshop on Computer Vision*, pages 79–85. IEEE Computer Society Press, 1987.
- [148] L. B. Wolff. Polarization-based material classification from specular reflection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(11):1059–1071, 1990.
- [149] L. B. Wolff. Constraining object features using a polarization reflection model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(6):635–657, 1991.
- [150] L. B. Wolff. Diffuse reflection. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society Press, June 1992.
- [151] R. J. Woodham. Photometric method for determining surface orientations from multiple views. *Optical Engineering*, 19(1):139–144, 1980.
- [152] R. Zhang, P.-S. Tsai, J. E. Cryer, and M. Shah. Analysis of shape from shading techniques. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 337–348. IEEE Computer Society Press, 1994.
- [153] Q. Zheng and R. Chellapa. Estimation of illuminant direction, albedo, and shape from shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):680–702, 1991.