Calibrated and Uncalibrated Photometric Stereo for Surface Texture Acquisition

Andrew D. Spence B.Sc. M.Sc.

Thesis submitted for the Degree of Doctor of Philosophy

Heriot-Watt University School of Mathematical and Computing Sciences



February 2005

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that the copyright rests with its author and that no quotation from the thesis and no information derived from it may be published without the prior consent of the author or the University (as may be appropriate).

Abstract

Pursuing a goal of realistic rendering for mixed reality applications using bump mapping demands the acquisition of both surface height and reflectance data of real textures. In this thesis we consider the use of various computer vision techniques for this purpose. We focus on establishing a practical implementation of Lambertian photometric stereo [Woodham1980]. Our objective is to make the technique more accessible so that it could potentially complement consumer-oriented visualisation applications. In this regard it is important to be unambiguous with respect to the standard operating procedure for optimal performance. It is also vital to minimise the requisite calibration of equipment.

With regard to three-image calibrated photometric stereo we determined that the optimal placement of the illumination vectors corresponds to an orthogonal arrangement. We also established that if the slant angle is constant, the optimal configuration corresponds to a difference of 120° between successive illumination tilt angles. Ignoring shadowing, we found the optimal slant angle to be a maximum of 90° for smooth surfaces and approximately 55° for rough surfaces.

With a view to reducing the requisite calibration, we developed a technique based on uncalibrated photometric stereo [Hayakawa1994]. It is practical to implement for surface texture planes and merely requires a single illumination tilt angle to be known. Its performance was found to be comparable to the equivalent calibrated technique. To my Father and in Memory of my Mother

Acknowledgement

I am greatly indebted to Dr. Mike Chantler for giving me the opportunity to return to university and work in a research environment. I must also thank him for the guidance, support and encouragement he has provided over the years not to mention coping with the recent proofreading burden. I'm grateful to Dr. Ged McGunnigle for his original work with the Texture Lab which established a welcome foundation for my research both in terms of theory and equipment.

I've very much enjoyed working in the Texture Lab and I would like to thank all of those people who've been there over the years for making it a great experience: Mike C, Ged, Cristina, Junyu, Jerry, Jennifer, Mike R, Kevin, Lingam, Ondrej, Ivan, Andreas, Michael and Maria. Thanks also to other friends and colleagues whom I've had the pleasure of meeting during my time at Heriot-Watt. Coffee breaks would have been very much duller without Jing and I always looked forward to the lunchtime powwow with Dave and Karina. Thanks also to Niall who introduced me to the wine appreciation society.

Thanks to friends and family for their unwavering support and encouragement. Finalmente, muchas gracias a Mariángeles por todo este año.

Table of Contents

Table	e of Con	itentsiv
List of Figuresviii		
List o	f Table	s xiii
Princ	ipal Sy	mbolsxiv
Abbr	eviation	1s xvi
Chap	ter 1	Introduction1
1.1	Motiv	ation1
1.2	Scope	
1.3	Criter	ia4
1.4	Contr	ibution5
1.5	Thesis	s Organisation
Chap	ter 2	Literature Review9
2.1	Introd	uction9
2.2	Reflec	ction of Light10
2.3	Mode	lling Light Reflection11
	2.3.1	Arbitrary Reflectance
	2.3.2	Reflectance Models
2.4	Religh	nting under Arbitrary Illumination15
	2.4.1	Criteria16
	2.4.2	Appearance-based Relighting16
	2.4.3	Model-based Relighting17
	2.4.4	Criteria Compliance
2.5	Lamb	ertian Photometric Stereo22
	2.5.1	Definitions22
	2.5.2	Assumptions & Limitations
	2.5.3	Mathematical Framework
	2.5.4	Summary
2.6	Sumn	nary & Discussion

Chapt	er 3	Surface Texture	31
3.1	Introd	uction	31
3.2	A Def	inition of Surface Texture	32
3.3	Surfac	e Description	33
3.4	Synthe	etic Surfaces	34
	3.4.1	Surface Models	35
	3.4.2	Data Generation	37
	3.4.3	Characterisation	38
3.5	Real V	Vorld Textures	41
	3.5.1	Data Generation	41
	3.5.2	Variable Colour Albedo	42
	3.5.3	Characterisation	46
3.6	Religh	ting Assessment	19
3.7	Summ	ary & Discussion	50

Chapt	er 4	Optimal Illumination for Three-Image Photometric Stereo	51
4.1	Introd	uction	51
4.2	Three	-Image Photometric Stereo	52
4.3	Noise	Expression Derivation	54
	4.3.1	Summary	59
4.4	Empir	ical Determination of Noise	60
4.5	Invest	igation into Optimal Performance	62
	4.5.1	Graphical Representation	62
	4.5.2	Minimisation	66
	4.5.3	Summary of Findings	69
4.6	Practi	cal Assessment	70
4.7	Summ	ary & Discussion	76
4.8	Concl	usions	77

Chapt	er 5	Uncalibr	rated Approach to Photometric Stereo	78
5.1	Introd	uction		78
5.2	Uncali	ibrated Ph	otometric Stereo	79
	5.2.1	Ambigui	ty Reduction/Resolution Survey	79
		5.2.1.1	Vector Magnitude Constraint	80
		5.2.1.2	Integrability Constraint	.82
		5.2.1.3	Consistent Viewpoint Constraint	85
5.3	Summ	ary & Dis	cussion	.89
5.4	Conclu	usion		.90

6.3 Summary & Discussion.....104

Chapter 7 Assessment of Uncalibrated Photometric Stereo

		with Texture - Simulation	106
7.1	Introd	uction	106
7.2	Exper	imental Technique Variations	107
7.3	Techn	ique Performance	
	7.3.1	Shadowing	109
	7.3.2	Specularities	113
	7.3.3	Point Illumination	118
	7.3.4	Presence of Noise	
	7.3.5	Illuminant Number & Positioning	126
	7.3.6	Mean Surface Normal Deviation	128
7.4	Summ	nary & Discussion	129

Chapt	er 8	Assessment of Uncalibrated Photometric Stereo	
		with Texture – Real Data	131
8.1	Introd	uction	131
8.2	Experi	imental Technique Variations	132
8.3	Assess	sment of Technique Performance	133
	8.3.1	Illuminant Number & Positioning	133
	8.3.2	Texture Character	136
	8.3.3	Equipment Set-up	146
8.4	Summ	ary & Discussion	148

Chapt	er 9 Summary and Conclusion	149
9.1	Summary	149
9.2	Conclusion	152
9.3	Contribution Recap	153
9.4	Future Work	

Appendix A: Texture-specific Uncalibrated Algorithm (UPS-tx) Results	155
Appendix B: Sensitivity Expressions for Common Slant Angle	186
Appendix C: Statistics Proofs	
Appendix D: Optimal Performance	
Appendix E: List of Publications by the Author	191
References	194

List of Figures

Figure 1.1	The appearance of a textile illuminated from different directions1
Figure 1.2	Utah teapot rendered in real-time using a photometrically-acquired textile bump map and texture map. Generated using an application written by Mike Robb
Figure 2.1	Reflection of light from a surface (a) specular, (b) glossy, (c) diffuse or matte
Figure 2.2	Diagram to illustrate the radiance definition11
Figure 2.3	Definition of BRDF parameters12
Figure 2.4	Illustration of vector geometry at surface facet for various reflectance models (a) Lambert, (b) Phong, (c) Torrance-Sparrow14
Figure 2.5	Equipment set-up for photometric stereo image capture22
Figure 2.6	Geometry of surface facet
Figure 3.1	The two components of the surface texture definition (a) colour albedo and (b) surface relief
Figure 3.2	PSD plot corresponding to (a) fractal & (b) sand ripple surfaces36
Figure 3.3	Height map images corresponding to (a) fractal, (b) Mulvaney, (c) sand ripple surfaces
Figure 3.4	Images of the (a) fractal, (b) Mulvaney, (c) sand ripple surfaces illuminated from two directions using the Lambertian model
Figure 3.5	Fourier spectra of images of the (a) fractal, (b) Mulvaney, (c) sand ripple surfaces illuminated from two directions $\tau=0^{\circ}$, 90°
Figure 3.6	Polar plot of (a) an isotropic texture and (b) a directional texture illuminated from directions with different tilt angles
Figure 3.7	Illustration of Bayer array for the top left corner of a camera greyscale image format
Figure 3.8	Image planes obtained by bilinear interpolation which are smaller in height & width due to form of the calculations43
Figure 3.9	One image from each of the real world texture data sets utilised in the investigation (illumination direction corresponds to $\tau=0^\circ, \sigma=45^\circ$)
Figure 3.10	Example images of the Fourier spectra of real world textures which exhibit (a) isotropic, (b) uni-directional and (c) multi-directional behaviour. These correspond to textures u, w and f respectively48
Figure 4.1	Equipment arrangement for three-image Lambertian photometric stereo
Figure 4.2	Reflectance maps for a Lambertian surface with illumination tilt angle of 0° and varying slant angle53
Figure 4.3	Flowchart depicting empirical sensitivity analysis60

Figure 4.4	Example of range of illumination configurations for two tilt angle experiments (Plan view). Increments are $\Delta \tau_3 = 1^\circ$ (theoretical), $\Delta \tau_3 = 10^\circ$ (empirical)
Figure 4.5	Figure of merit M_{rough} versus third tilt angle τ_3 with $\tau_1=0^\circ$, $\tau_2=120^\circ$
Figure 4.6	Figure of merit M_{rough} versus third tilt angle τ_3 with $\tau_1=0^\circ, \tau_2=90^\circ63$
Figure 4.7	Illustration of range of illumination slant angles in experiments. Increments are $\Delta \sigma = 1^{\circ}$ (theoretical), $\Delta \sigma = 10^{\circ}$ (empirical)64
Figure 4.8	Total noise ratio M_{rough} versus slant angle σ with $\tau_1=0^\circ$, $\tau_2=120^\circ$, $\tau_3=240^\circ$
Figure 4.9	Total noise ratio M_{smooth} versus slant angle σ with $\tau_1=0^\circ$, $\tau_2=120^\circ$, $\tau_3=240^\circ$
Figure 4.10	M_{rough} versus $\mathbf{b}_1.\mathbf{b}_2$ and $\mathbf{b}_2.\mathbf{b}_3$
Figure 4.11	Representation of orthogonal vectors
Figure 4.12	Practical evaluation of the proposed optimal illumination conditions71
Figure 4.13	Comparison of original image with relights72
Figure 4.14	Moving average (10-pixel) trend for a single row of the difference images given in Figure 4.1274
Figure 4.15	Practical evaluation of effect of slant angle on accuracy based on optimal illumination configuration with regard to tilt angle $(\Delta \tau = 120^{\circ})$
Figure 5.1	Mirror-like specular reflection from a surface facet such that the highlight is visible
Figure 6.1	Flowchart of uncalibrated photometric stereo first stage which entails formulation of the intensity matrix & its decomposition into pseudo surface & illumination matrices
Figure 6.2	Flowchart of uncalibrated photometric stereo second stage in which the ambiguity is determined in order to transform the pseudo surface & illumination matrices into actual estimates
Figure 6.3	Flowchart of initial ambiguity reduction in texture-specific photometric stereo
Figure 6.4	Flowchart of ideal residual ambiguity resolution in texture-specific photometric stereo
Figure 6.5	Surface normal distribution for synthetic fractal texture corresponding to pseudo matrix estimate (a) before & (b) after <i>z</i> -axis alignment
Figure 6.6	Pseudo illumination vector corresponding to known vector requires rotation of α_2 to orient <i>x</i> - <i>y</i> plane with viewer-oriented co-ordinate system
Figure 6.7	Flowchart of practical residual ambiguity resolution in texture-specific photometric stereo103

Figure 7.1	Visual explanation of both self and cast shadowing109
Figure 7.2	Sample height map profile for fractal surface110
Figure 7.3	Sand ripple surface of increasing surface roughness illuminated under identical conditions
Figure 7.4	Mean <i>TSER</i> for three synthetic textures versus rms roughness for the UPS-tx algorithm
Figure 7.5	Sample output images for sand ripple surface (rms roughness=5)112
Figure 7.6	Mean <i>TSER</i> for three synthetic textures versus rms roughness for five Lambertian photometric stereo algorithms
Figure 7.7	Reflection from a surface facet (a) diffuse reflection modelled by Lambert's law, (b) specular reflection modelled by the Phong model with an exponent of 10 and (c) specular reflection modelled by the Phong model with an exponent of 100113
Figure 7.8	Mean <i>TSER</i> for three synthetic textures versus proportion of specular reflection over a range of Phong exponent values for the UPS-tx algorithm
Figure 7.9	Mean <i>TSER</i> for three synthetic textures versus Phong exponent over a range of specular reflection proportions for the UPS-tx algorithm
Figure 7.10	Phong reflection with exponent of (a) 2, (b) 10 and (c) 100. Graphs generated from spreadsheet created by Dr. M. J. Chantler115
Figure 7.11	Scaled images of a Mulvaney surface exhibiting (a) diffuse reflection, (b) specular reflection (Phong exponent = 10) and (c) specular reflection (Phong exponent = 100)116
Figure 7.12	Sample output images for specularly reflecting Mulvaney surface data set corresponding to a Phong exponent of 10 and a specular reflection proportion of 50%116
Figure 7.13	Mean <i>TSER</i> for three synthetic textures versus proportion of specular reflection for four Lambertian photometric stereo techniques (Phong exponent=40)117
Figure 7.14	Mean <i>TSER</i> for three synthetic textures versus Phong exponent for four Lambertian photometric stereo techniques (specular reflection proportion=0.4)
Figure 7.15	Point illumination diagram illustrating difference in illumination vectors with regard to both magnitude and direction118
Figure 7.16	Point illumination calculations (a) light source co-ordinate, (b) illumination angles and radius for a facet
Figure 7.17	Mean <i>TSER</i> for three synthetic textures versus illumination radius for square sample of side (a) 0.4m & (b) 0.1m for five photometric stereo algorithms

Figure 7.18	Images of fractal surfaces illuminated by a nearby point light source with illumination radius between 0.5m and 2.0m (a) images not scaled with each other to facilitate illumination gradient evaluation (b) images scaled with each other
Figure 7.19	Sample output images for data set corresponding to square fractal surface of side 0.4m under point illumination of radius 0.5m122
Figure 7.20	Images of (a) white noise, (b) illuminated fractal surface and (c) image resulting from the addition of (a) and (b). Signal to noise ratio is ~10dB124
Figure 7.21	<i>TSER</i> versus input signal to noise ratio for an illuminated fractal surface for four Lambertian photometric stereo techniques
Figure 7.22	Sample output images for the fractal data set corresponding to an input signal to noise ratio of 10dB
Figure 7.23	Mean <i>TSER</i> for three synthetic surfaces versus number of images input to three photometric stereo techniques. Illuminant configurations largely asymmetric
Figure 7.24	Mean <i>TSER</i> for three synthetic textures versus number of images input to three photometric stereo techniques. Illuminant configurations symmetric
Figure 7.25	Mean <i>TSER</i> for three synthetic textures versus mean surface normal offset angle for the UPS-tx algorithm
Figure 8.1	Mean <i>TSER</i> versus number of images input to three photometric stereo techniques. Mean <i>TSER</i> averaged over 31 real textures. Illuminant configurations largely asymmetric
Figure 8.2	Mean <i>TSER</i> versus number of images input to three photometric stereo techniques. Mean <i>TSER</i> averaged over 31 real textures. Illuminant configurations symmetric
Figure 8.3	<i>TSER</i> (texture signal to relight error ratio) for real textures with input images which are (a) symmetrically lit and (b) asymmetrically lit
Figure 8.4	Relative performance ranking of the real textures for the UPS-tx algorithm corresponding to a symmetric illumination configuration
Figure 8.5	Surface representations of texture <i>g</i> generated by the UPS-tx algorithm
Figure 8.6	Intensity profile for a specific pixel position (number 1000) in a series of 36 16-bit images of texture <i>g</i> corresponding to changing illumination direction of monotonically increasing tilt angle and constant slant angle. Diffuse fit determined using Saito's method [Saito1996]141
Figure 8.7	Intensity images of texture <i>ac</i> with specular highlights (a) not visible and (b) visible

Figure 8.8	Intensity profile for a specific pixel position (number 1000) in a series of 36 16-bit images of texture <i>ac</i> under changing illumination direction of monotonically increasing tilt angle and constant slant angle. Diffuse fit and Phong fit determined using Saito's method [Saito1996]143
Figure 8.9	Original image of texture <i>ae</i> (τ =90°, σ =60°), a relit image (τ =90°, σ =60°) generated by using surface representations estimated by UPS-tx with a symmetric illumination configuration & their difference image
Figure 8.10	Intensity profile for a pixel position (corresponding to the cement section) in a series of 16-bit 36 images of texture <i>ae</i> under changing illumination direction of monotonically increasing tilt angle and constant slant angle
Figure 8.11	Original image of texture h (τ =230°, σ =45°), a relit image (τ =230°, σ =45°) generated by using surface representations estimated by UPS-tx with a symmetric illumination configuration & their difference image
Figure 8.12	Intensity profile for a pixel position (corresponding to non-Lambertian reflection) in a series of 36 16-bit images of texture <i>h</i> under changing illumination direction of monotonically increasing tilt angle and constant slant angle. Diffuse fit determined using Saito's method [Saito1996]146

List of Tables

Table 3.1	Measures of surface structure features for the three synthetic surfaces	38
Table 3.2	Measures of surface structure features for the real world surface textures	16
Table 4.1	Examples of three parameter minimisation results	66
Table 4.2	Examples of four parameter minimisation results for (a) a rough surface, (b) a smooth surface	67
Table 4.3	Examples of six parameter minimisation results for a rough surface	58
Table 5.1	Summary of literature review of uncalibrated photometric stereo with emphasis on ambiguity reduction/resolution	39
Table 8.1	Relative performance ranking of the real textures for the photometric stereo algorithms PS, UPS-tx and UPS with either symmetric or asymmetric illumination configuration	.39
Table 8.2	Mean TSER (texture signal to relight error ratio) averaged over thirty-one real textures for different sizes of texture sample	147

Principal Symbols

Symbol	Meaning	Introduced in Section:
Α	Surface area (m^2)	2.3
Α	Ambiguity matrix	2.5
а	Element of A	5.2
В	Illumination unit vector matrix	2.5
b	Illumination direction unit vector	2.3
С	Symmetric matrix	5.2
с	Element of C	5.2
d	Directionality measure	3.3
E_i	Irradiance, the flux incident on a surface (Wm ⁻²)	2.3
F	Fresnel reflectivity	2.3
f	Number of frames i.e. number of input images corresponding independent illumination directions.	to 2.5
fr	Bidirectional reflectance distribution function	2.3
G	Matrix	5.2
g	Vector	5.2
ь Н	Element of H	5.2
h	Vector composed of elements of H_1 and H_2	5.2
H	Skew symmetric matrix	5.2
I	Image intensity image or matrix	2.5
Ī.	Identity matrix	5.2
- <i>a</i> i	Image intensity vector	2.5
i	Pixel intensity value & element of I or i . (Subscript <i>L</i> , <i>P</i> , <i>TS</i> ref Lambertian, Phong and Torrance-Sparrow models)	er to 2.5
k	Constant (Subscript <i>f,m,s</i> for fractal, Mulvaney, sand ripple P subscript <i>ov</i> in surface normal estimates)	SD; 3.4
L	Scaled illumination vector matrix	2.5
ŕ	Pseudo scaled illumination vector matrix	2.5
	Scaled illumination vector	2.5
î	Pseudo scaled illumination vector	2.3 5.2
l î	Pseudo scaled illumination vector element of $\hat{\mathbf{L}}$ and $\hat{\mathbf{l}}$	6.2
L _r	Radiance, the flux emitted from a surface into a solid angle ce around the direction of the radiated light $(Wm^{-2}sr^{-1})$.	entred 2.3
т	Number of pixels in an image	2.5
n	Surface normal unit vector	2.3
n	Phong exponent	2.3
р	Surface facet gradient in the x-direction.	2.5
p_{rms}, q_{rms}	Root mean square slope	3.3
Q	Geometric attenuation factor	2.3
q	Surface facet gradient in the y-direction.	2.5
R	Residual ambiguity matrix	5.2
Rz	Rotation about <i>z</i> -axis matrix	5.2
R	Real number	5.2
<i>r</i> _b	Illumination radius i.e. distance of light source from texture (m) 7.3
r	Tangent to surface facet (Subsript x, y denote x-, y- direction)	2.5
S	Scaled surface normal matrix	2.5
Ŝ	Pseudo scaled surface normal matrix	2.5

S	Scaled surface normal	2.5
S	Scaled surface normal element of S or s	5.2
S	Power spectral density, PSD	3.4
	(Subscript <i>f,m,s</i> for fractal, Mulvaney, sand ripple surfaces)	
ŝ	Pseudo scaled surface normal	5.2
ŝ	Pseudo scaled surface normal element of $\hat{\mathbf{S}}$ and $\hat{\mathbf{s}}$	5.2
SER	Signal to relight error ratio (dB)	3.6
TSER	Texture signal to relight error ratio (dB)	3.6
U	Orthogonal matrix produced by singular value decomposition	2.5
(u,v)	Cartesian frequency coordinates	3.4
V	Orthogonal matrix produced by singular value decomposition	2.5
v	Viewing direction unit vector	2.3
W	Pseudo viewing vector	5.2
(x,y)	Pixel location in image	2.3
Z	Surface facet height	2.5
Zrms	Root mean square (rms) roughness	3.3
α	Albedo term, the product of albedo and light source intensity	2.5
β	Roll-off factor	3.4
θ'	Bisecting angle between the lighting and viewing directions	2.3
$ heta_a$	Angle between the surface normal and the bisector of the illumination	2.3
	and viewing directions (r)	
$ heta_i$	Angle between surface normal and illumination vector (r)	2.3
θ_r	Angle between surface normal and viewing vector (r)	2.3
λ	Light source intensity	2.3
ρ	Diffuse albedo	2.3
ρ_s	Specular albedo	2.3
σ	Slant angle, the angle between the illumination vector and the z-axis	2.5
$\boldsymbol{\Sigma}$	Diagonal matrix produced by singular value decomposition	2.5
τ	Tilt angle, the angle between the x-axis and the projection of the	2.5
	illumination vector onto the x-y plane	
Φ	Flux, the rate of incident/emitted light energy (W)	2.3
Ψ	Noise/standard deviation.	4.3
ω	Radial frequency	3.4
Ω_r	Solid angle of reflected beam (sr)	2.3
Ω_r	Solid angle of incident beam (sr)	2.3
ϕ_r	Angle to projected viewing vector in the plane (r)	2.3
ф.	Angle to projected illumination vector in the plane (r)	2.3
T l		

Abbreviations

Abbreviatio	on Meaning	Introduced in Section:
3D/2D	Three/two-dimensional	1.2
ACF	Autocorrelation function	3.3
BRDF	Bidirectional reflectance distribution function	2.3
BTF	Bidirectional texture function	2.3
CCD	Charge coupled device	1.2
CUReT	Columbia Utrecht Reflectance and Texture database	2.3
GBR	Generalized bas-relief transformation	5.2
GL(3)	General linear transformation	5.2
O(3)	Orthogonal transformation	5.2
PC	Personal Computer	1.3
PS	Lambertian photometric stereo algorithm (>3 input images)	7.2
PS-3	3-image Lambertian photometric stereo algorithm	7.2
PSD	Power spectral density function	3.3
PTM	Polynomial texture map	2.4
RGB	Red, green, blue	3.5
rms	Root mean square	3.3
SFS	Shape from Shading	2.4
SVD	Singular Value Decomposition	2.5
TIFF	Image format	3.5
T-S	Torrance-Sparrow reflectance model	2.4
UPS	Uncalibrated photometric stereo algorithm	7.2
UPS-it	Iterative UPS-tx algorithm	7.2
UPS-tx	Texture-specific uncalibrated photometric stereo algorithm	7.2

Chapter 1

Introduction

1.1 Motivation

This thesis is concerned with the acquisition of three-dimensional surface texture data and corresponding surface reflectance data. The motivation for this work is to enable more accurate modelling of rough patterned surface textures in mixed reality visualisation applications. This technology is relevant to areas such as advertising, entertainment and computer-aided design where photorealistic rendering is a desirable and potentially valuable tool.



Figure 1.1 The appearance of a textile illuminated from different directions.

Texture mapping is a computer graphics technique which is commonly used to enhance the realism of a virtual scene. It involves the use of an image of a texture to tile the surface of a polygon object. However, in this case only the shape of the object is taken into account with regard to scene illumination. Since the appearance of rough surface textures depends on both viewpoint and illumination direction [Chantler1995], it is important to model this effect to attain more convincing levels of realism (see Figure 1.1). Bump mapping [Blinn1978] is a computer graphics technique which can be used to represent the surface relief associated with texture. Used in conjunction with texture mapping it enables surface texture to be included in lighting calculations. Until recently its implementation on standard PC equipment was limited to software-based calculation and scene rendering times were of the order of seconds. With the advent of inexpensive programmable PC graphics cards, however, rendering in real-time is now feasible from an economic point of view. This is due to the fact that both bump mapping and texture mapping techniques are compatible with the cards and can therefore be implemented in hardware [Robb2003]. As a result this technology is likely to find wider use in practical visualisation applications.



Figure 1.2 Utah teapot rendered in real-time using a photometrically-acquired textile bump map and texture map. Generated using an application written by Mike Robb. Scene updated on change of illumination direction, object pose or viewing position (70 frames per second).

Pursuing a goal of realistic rendering in this way demands generating the bump maps of real textures. Photometric stereo [Woodham1980] is one computer vision technique which provides a means of determining the requisite data. It allows the surface normals from which the bump map is derived and the corresponding reflectance which provides a colour texture map to be estimated. The focus of this work is to investigate methods for obtaining this data with a view to establishing a practical implementation which could potentially complement consumer-oriented visualisation applications. The intention is to make the selected techniques more accessible so that for example, a product designer could independently create a visualisation based on selected textile samples taken from a swatch using a simple image acquisition system. In this regard it is important to be unambiguous with respect to the standard operating procedure for optimal performance. It is also vital to minimise the requisite calibration of equipment. It is these issues in particular which are considered in this thesis.

Three-dimensional surface texture data and corresponding surface reflectance data not only facilitate visualisation applications but are employed for other purposes. This may mean that a practical implementation of a suitable technique finds further application in the field of computer vision. Other uses include segmentation and classification [McGunnigle1998, McGunnigle2001], surface height recovery for inspection [Smith1999a, Smith1999b, Gullón2002], texture synthesis [Dong2002, Dong2003a], face recognition [Georghiades2003a] and forensic science [McGunnigle2001].

1.2 Scope

Our objective is to facilitate the relighting of rough patterned surface textures by generating a suitable surface representation. A number of approaches have been developed in the field of computer vision with regard to recovering three-dimensional surface data. Contact methods, which generally employ a stylus which is moved across the surface, are outwith the scope of the thesis. We concentrate on non-contact methods although disregard the use of direct measurement methods such as time-of-flight due to the equipment expense and complexity. We assume the use of a single CCD camera in a fixed position and hence do not consider multiple view techniques such as stereo vision and optical flow. These methods are better suited to determining the shape of Furthermore they do not recover the reflectance objects rather than texture. characteristics of the object surface which is necessary for relighting. We focus on producing bump maps and corresponding albedo images by processing colour images of a surface texture captured from a fixed camera under different illumination directions using the shape from shading technique photometric stereo. Other representations which can be generated with image-based techniques are also discussed, however.

Whilst the surface representations can be utilised in the generation of 3D visualisations such as Figure 1.2, the specific details of the implementation for this, such as the design of the graphics pipeline, are not covered in this thesis. When we refer to relighting, we mean that an image of the texture is generated under a novel lighting direction. We use these images to assess the performance of the various algorithms tested. As mentioned, the appearance of rough surface textures changes under both viewpoint and illumination

position. We consider only the latter in the numerical assessment and provide visual examples of the former when the bump map and albedo texture map are applied to an object.

1.3 Criteria

Although the scope of this work is limited to images of textures captured from the same viewpoint but under different illumination directions, there remain a range of techniques which can be employed to process the intensity data and hence generate suitable surface representations. In order to distinguish between them and identify those most suited to our needs we use the following criteria:

1. Suitable for globally flat diffuse surfaces

The texture should be globally flat such that its mean surface normal is aligned with the line of sight of the camera. Surface reflection should be able to be approximated by the Lambertian model.

2. Use of consumer-level equipment.

Image capture should be possible with inexpensive camera equipment. Specialised hardware should not be required for real-time relighting. A PC or laptop with a consumer-level programmable graphics card should be sufficient for this purpose.

3. Practical input data capture.

The image capture procedure should be simple, straightforward and not time-consuming. As few images as possible should be required. Calibration should be avoided where possible. Calibration objects should not be used.

4. Computationally efficient generation of representation.

The algorithm used to produce the representation should be 'fast'. Ideally iterative techniques should be avoided.

5. Compact and compatible representation.

The representation should be of a low dimension. It should also be compatible with consumer-level programmable graphics cards.

6. Computationally efficient generation of accurate relit images.

Real-time per-pixel rendering should be possible with the graphics card and the representation. The resulting images should be accurate and convincing. This means that the representation should be able to approximate the reflectance characteristics of the surface texture to a sufficiently high level.

To summarise: the ideal technique uses inexpensive equipment, requires few input images and minimal calibration. The resulting representation is fast to compute, compact i.e. of low dimension, compatible with programmable graphics cards and provides accurate relit images of diffuse surface textures.

1.4 Contribution

This thesis makes the following contributions:

- 1. The optimal operating conditions are determined for the three-image Lambertian photometric stereo technique. Prior to this work, the use of the maximum practical illumination slant angle and the avoidance of co-planar illumination arrangements were the sole recommendations. We show that the relative arrangement of the three illumination vectors has an impact on performance. We prove that an orthogonal arrangement is optimal through the use of a derived figure of merit. On a practical basis whereby the slant angle is constant we prove and verify by experiment that a tilt angle separation of 120° is optimal. This helps to clarify the standard operating procedure for the method.
- 2. We propose the use of the figure of merit to gauge the effect of various illumination configurations on performance. We show that McGunnigle's simplified photometric stereo scheme [McGunnigle1998] which uses a difference in tilt angle of 90° is sub-optimal but not significantly so. We believe that this is also relevant to robust photometric stereo techniques which use more than three images but select the best three intensity values at each pixel position. Discarding the outlying intensity values may mean that the illumination configuration corresponding to the remaining intensities is far from optimal.
- 3. We identify an uncalibrated version of photometric stereo as a promising candidate for bump map and albedo image estimation with minimum equipment calibration. A practical implementation of the method is developed for specific use with surface texture planes. For relighting purposes, it assumes that the light source intensity is constant and that a single illumination tilt angle is known. The technique is tested in simulation with synthetic textures and with thirty-one real textures. We thus demonstrate that it is of practical use and attains accuracy levels comparable to the equivalent calibrated technique.

1.5 Thesis Organisation

This thesis is organised into nine chapters. In Chapter 2 we present a review of relevant literature and we select several approaches which best satisfy the criteria stated in Section 1.3. The synthetic and real textures whose images are used as input data to test the performance of the selected methods are described in Chapter 3. In Chapter 4 a sensitivity analysis of one of the selected methods, three-image Lambertian calibrated photometric stereo, is described with a view to determining the optimal illumination arrangement. The uncalibrated photometric stereo method is detailed in Chapter 5 and subsequently developed for texture-specific use in Chapter 6. The results of a series of simulation experiments with both calibrated and uncalibrated techniques are reported in Chapter 7. In Chapter 8 their performance with real textures is reported. Finally we summarise the thesis and draw conclusions in Chapter 9.

A brief synopsis of each chapter is given in the following:

Chapter 2

We initially consider the processes which take place when light is reflected from a surface and review the models which are used to describe them. We then survey the various approaches taken in the literature with regard to relighting a textured surface. We determine that these fall into two groups: model-based and image-based methods. For the purposes of this thesis we establish that the methods most suited to our needs are model-based shape from shading algorithms by considering the stated criteria. In particular we single out photometric stereo in both calibrated and uncalibrated forms for further investigation.

Chapter 3

We define surface texture as rough surfaces with a planar megastructure which are of colour and potentially patterned. We introduce three synthetic and thirty-one real textures which satisfy this definition. The images of these textures, both generated and captured, form the input data which are to be processed by the selected algorithms. We characterise the textures in terms of surface roughness measures and we also consider their second order statistics.

Chapter 4

We outline the approach taken to evaluate the optimal illumination arrangement for three-image calibrated photometric stereo. We describe a sensitivity analysis which is utilised to examine the presence of noise in the output images. We derive an overall figure of merit via this approach. It is an equation which approximates the total noise in a relit image and is expressed in terms of the illumination tilt and slant angles. We prove the optimal illuminant arrangement by using the figure of merit and verify the result through experiment by analysing the data for each of the thirty-one real textures.

Chapter 5

We introduce the uncalibrated photometric stereo method in detail and indicate that the main issue involved concerns the reduction and resolution of the inherent ambiguity. We consider the various methods employed to solve it. These are based on three main constraints. We reject the use of the integrability constraint since we are working with rough potentially discontinuous surfaces. Furthermore, we reject the use of the consistent viewpoint constraint since this requires a specular reflecting surface and the majority of the thirty-one real textures exhibit diffuse reflection. We point out that the constant light source intensity which constraints the magnitude of the illumination vectors to be constant is a potentially suitable approach.

Chapter 6

We consider the constant light source intensity approach to ambiguity reduction for uncalibrated photometric stereo in greater detail. We determine that the residual ambiguity is orthogonal in nature and consider ways to resolve it without resorting to complete calibration. We develop a stepwise method which uses the fact that the textures are globally flat to orient the vector field with regard to the *z*-axis. For relighting purposes, a single illumination tilt angle is merely required to provide a rotation about the *z*-axis.

Chapter 7

We test a number of variants of photometric stereo, both calibrated and uncalibrated methods, under various conditions. Since Lambertian photometric stereo assumes ideal diffuse reflection, we evaluate the performance of the algorithms with shadowing, specularities and point lighting. We also test the algorithms with noisy input. We investigate the effect of varying the number of input images and the relative arrangement of the corresponding illumination directions.

Chapter 8

We test the proposed texture-specific method with the thirty-one real textures and compare its performance against that attained by the other techniques. We investigate the effect of varying the number of input images and the relative arrangement of the corresponding illumination directions. We also examine the general impact of texture character on relighting accuracy.

Chapter 9

We summarise the thesis and draw conclusions. We indicate that we have determined the optimal operating conditions with regard to the illumination configuation for 3-image photometric stereo. We propose the uncalibrated photometric stereo algorithm we developed as a practical method for the capture of bump maps and albedo images of surface texture.

Chapter 2

Literature Review

2.1 Introduction

When light is incident on a surface, it may be reflected, absorbed and transmitted. The extent to which each process occurs depends on the material. This page provides a physical illustration: the ink of the printed letters absorbs the light it encounters and thus appears black whilst the white paper reflects it. Transmission also occurs through the paper and can be observed if the single sheet of A4 is held up to the light. With regard to the appearance of illuminated surface textures, it is only the light which reaches the viewer from the surface which we will consider. In this thesis we are concerned with the *reflectance* of light. In this chapter we briefly review the physical processes involved in order to appreciate the rationale behind the various reflection models described. The formulation of these models suggests one approach for relighting under arbitrary illumination which is based on obtaining representations of both surface geometry and reflectance. We present a literature survey of existing techniques which determine such representations from intensity images of a surface. Alternative approaches which do not explicitly determine surface orientation are also detailed. The overall goal of this chapter is to identify relighting techniques from the literature which satisfy the criteria specified in Section 1.3.

This chapter is organised as follows :

The reflection of light and ways of modelling it are briefly introduced in Sections 2.2 and 2.3. A comprehensive but concise survey of approaches to relighting is presented in Section 2.4. The mathematical procedure for implementing some of the suitable methods is detailed in Section 2.5. We summarise the chapter in Section 2.6.

2.2 Reflection of Light

Two main processes are often assumed to account for the total light reflected from an object and are briefly considered here.

Interface reflection

If the object material is homogeneous in an optical sense then it has a uniform refractive index throughout. The implication is that a ray of light will not be able to penetrate the body of the object and will effectively be reflected from the air-object interface. If the object surface is flat then the reflected ray forms a single well-defined beam which mirrors the corresponding incident ray (Figure 2.1a). This is known as *specular* reflection. It is noted that the reflected light retains the original spectral characteristics of the incident light because it effectively does not interact with the material of the object.

With regard to surface roughness, if the surface irregularities are relative in size to the wavelength of the incident light, the resulting light beams will be scattered to some extent. Interface reflections still take place but the varied surface orientation ensures a variety of reflection directions (Figure 2.1b). If the deviations in surface height are small, the object will appear *glossy* rather than specular. If the deviations are large the object will appear *matte*. Interface reflection from a very rough surface therefore contributes to *diffuse reflection*.



Figure 2.1 Reflection of light from a surface (a) specular, (b) glossy, (c) diffuse or matte.

Body scattering

A common mechanism for diffuse reflection is body scattering. In this case the object material is inhomogeneous and the light penetrates the surface due to refraction at the surface. It subsequently undergoes refraction and reflection at interfaces between regions of differing refractive indices in the body of the object thus scattering the light internally. Some eventually reaches the surface and radiates back into the air in random directions (Figure 2.1c). During this subsurface scattering process interactions with the object material occur and result in changes in the spectral characteristics of the reflected

light. The resulting colour depends on both the reflectance spectrum of the object material and the spectrum of the incident light in this case.

2.3 Modelling Light Reflection

Before considering the various approaches taken to model the reflectance of light from a surface it is useful to state the standard radiometric terms which will be used in their development. *Flux* Φ is the rate of incident/emitted light energy measured in watts (W). *Irradiance* $E_i = d\Phi/dA$ and is the flux incident on a surface measured in watts per unit surface area (Wm⁻²). *Radiance* $L_r = d\Phi/(dA \cos\theta_r d\Omega_r)$ and is the flux emitted from a surface into a solid angle¹ centred around the direction of the radiated light measured in watts per unit foreshortened area per steradian (Wm⁻²sr⁻¹). See Figure 2.2 for an illustration. The foreshortened area is the area of illuminated surface patch multiplied by the cosine of the angle between the radiated light direction **v** towards the viewer and the surface normal **n**. It accounts for the effective surface area seen from the standpoint of the viewer.



Figure 2.2 Diagram to illustrate the radiance definition.

It is important to note that the radiance term is related to the brightness of the surface patch observed by the viewer i.e. the pixel intensity value in an image of the surface.

¹ For a cone it is the ratio of the area 'cut out' of a sphere with its centre at the cone apex to the square of the sphere's radius.

2.3.1 Arbitrary Reflectance

Surface reflectance can be accurately described by the bidirectional reflectance distribution function (BRDF) [Nicodemus1977]. In standard radiometric terms it is defined as the ratio of the reflected radiance to incident irradiance and therefore takes units of inverse steradians. The BRDF is usually expressed as a four parameter function:

$$f_r(\theta_i, \phi_i, \theta_r, \phi_r) = \frac{dL_r(\theta_i, \phi_i, \theta_r, \phi_r)}{dE_i(\theta_i, \phi_i)}$$
(2.1)

The diagram in Figure 2.3 depicts the relevant angles. This version of the BRDF assumes monochromatic light and ignores subsurface scattering such that the outgoing position is the same as the incoming. The BRDF represents a measure of the brightness of the surface viewed from one direction when it is illuminated from another direction. The reflectance is determined by integrating the function over a range of incident and reflected angles.



Figure 2.3 Definition of BRDF parameters.

The BRDF can hence be used to optically characterise the material of an object. Real world materials are likely to be inhomogeneous, however, since both reflectance and microgeometry may vary on a local basis. The bidirectional texture function (BTF) was defined to account for this [Dana1999]. The BTF extends the BRDF by allowing it to vary spatially across the surface to capture its *texture*. It is hence a six parameter function. The distinction between the BRDF and the BTF depends on the scale of observation.

Measuring either the BRDF or the BTF is challenging due to the high dimensionality of the data. With regard to the BRDF, robot arms have been utilised to precisely place a photometer and light source over the hemisphere of potential directions above a planar sample of material; a single measurement is taken at each position. This is time-consuming even for a sparse sampling of the function and ways to improve the efficiency have been presented e.g.[Ward1992, Lu1998, Lu2000, Marschner1998]. Dana *et al* used a robotic system to sample the BTF for 61 textures collecting over 200 images for each and these form the CUReT database. A novel technique for measuring the BTF *in situ* with no mechanical movement was recently presented whereby a kaleidoscope enables the surface to be viewed simultaneously from many directions [Han2003]. The acquisition of densely sampled data sets for the BTF of arbitrary surfaces was recently reported; each texture data set contains around *10,000 images* [Koudelka2003].

2.3.2 Reflectance Models

Reflectance models are commonly used in both computer graphics and computer vision as a practical alternative to the empirical BRDF & BTF. Their relative simplicity is advantageous although the range of applicability for those with fewer parameters is restricted.

Diffuse Reflection

The oldest model was presented by Lambert in 1760. It describes reflectance from a perfectly diffuse surface. In this case the appearance of the illuminated surface is assumed to be the same (i.e. equally bright) from all viewing directions. The implication is that the Lambertian BRDF is *constant*. The brightness is proportional to the cosine of the angle between the illumination vector and the surface normal θ_i (see Figure 2.4a). Although empirically-derived, this effect can be attributed to the foreshortened surface area from the point of view of the light source. Ignoring shadowing/interreflections and assuming a point light source at infinity, an image pixel intensity for a Lambertian surface is given as follows:

$$i_{L}(x, y) = \rho \lambda \cos \theta_{i} = \rho \lambda \mathbf{n} \cdot \mathbf{b}$$
(2.2)

where ρ is the diffuse albedo and is defined as the proportion of incident light which is reflected as diffuse light, λ is the light source intensity.

Oren and Nayar reported deviations from ideal diffuse reflection for rough surfaces [Oren1994]. They generalised Lambert's model to take account of the increase in brightness observed as viewing direction approaches the lighting direction. Other contributions to the development of more accurate physically-based diffuse reflection

models include that for smooth surfaces [Wolff1994], layered materials [Hanrahan1993] and both rough and smooth surfaces [Wolff1998]. Despite such improvements, Lambert remains popular because this simple model gives reasonable results for a wide range of matte surfaces. Furthermore, it often features as the diffuse component in hybrid models of both body scattering and interface reflection.



Figure 2.4 Illustration of vector geometry at surface facet for various reflectance models (a) Lambert, (b) Phong and (c) Torrance-Sparrow.

Specular Reflection

The BRDF of a perfectly specular reflecting surface can be modelled as a Dirac delta function [Horn1989, Chap.8]. This is equivalent to the case of interface reflection from a flat surface such that the reflected light direction mirrors that incident about the surface normal as illustrated in Figure 2.1a. In contrast to the perfect diffuse model which is often adequate for describing reflection from real matte surfaces, this perfect specular model will not produce satisfactory results for real world surfaces which exhibit specular highlights. Parametric models have been developed for such non-Lambertian reflectance. In general these are *hybrid* models and are a linear combination of a diffuse and a specular component. It is noted that an ambient component may also be included but as this merely acts as offset it will not be considered here.

The most commonly-used parametric model for specular reflection in computer graphics applications is the empirically-based Phong model [Phong1975]. It is a linear combination of Lambertian diffuse reflection and a cosine function raised to a power which corresponds to a specular lobe. An image pixel intensity for a specularly reflecting surface is given as follows by the Phong model:

$$i_P(x, y) = i_L(x, y) + \lambda \rho_s (\mathbf{v} \cdot \mathbf{r})^n$$
(2.3)

where **v** is the viewing vector, **r** is the mirror reflection of the incident light vector, ρ_s is the specular albedo (the fraction of incident light reflected as specular light) and *n* is the Phong exponent (See Figure 2.4b). The vector \mathbf{r} can be written in terms of the surface normal and the illumination vector such that $\mathbf{r} = 2(\mathbf{b} \cdot \mathbf{n})\mathbf{n} - \mathbf{b}$.

Although it provides reasonable results the Phong model has no physical basis. For example, it does not model the off-specular peak observed for both many metallic and non-metallic surfaces. The Torrance-Sparrow model (T-S) is derived from geometrical optics and does model this effect [Torrance1967]. An image pixel intensity for a specularly reflecting surface is given as follows by the T-S model:

$$i_{TS}(x,y) = i_L(x,y) + \lambda \rho_s Q F(\theta',\eta) \frac{\exp(-\upsilon^2 \theta_a^2)}{\cos \theta_r}$$
(2.4)

where v is the surface roughness, θ_a is the angle between the surface normal and the bisector of the illumination and viewing directions, Q is the geometric attenuation factor to account of masking and shadowing, and F is Fresnel reflectivity which depends on the refractive index and the bisecting angle θ' between the lighting and viewing directions (see Figure 2.4c).

Various improvements to the T-S model have been suggested. Blinn augmented the T-S model with an alternative distribution function [Blinn1977]. Nayar *et al* proposed a hybrid model which unites the geometric optics and physical optics approaches. Their model consists of a diffuse lobe, a specular lobe and a specular spike [Nayar1991]. These components are based on Lambert, Torrance-Sparrow and Beckmann-Spizzichino respectively. Cook and Torrance proposed a reflectance model which takes into account the spectral appearance [Cook1982]. The dichromatic reflectance model is based on an interface reflection component and a body scattering component which are each described with a geometrical and a spectral term [Shafer1985]. He *et* al extended the Cook-Torrance model to include specular reflection for reduced surface roughness and introduced a directional diffuse term [He1991].

2.4 Relighting under Arbitrary Illumination

In the previous section we considered various ways to either represent or model the reflection of light from surfaces with different reflectance characteristics. It is readily apparent that if the parameters of a model are known as is the surface geometry then it is possible to generate image intensities which correspond to an arbitrary illumination direction i.e. by substituting the corresponding vector into the equation. In this section we survey the techniques which can be utilised to simultaneously determine geometry and model parameters to facilitate arbitrary relighting in this manner. We do not

consider separate determination of surface geometry through the use of range finders e.g. [Ikeuchi1991, Sato1997], or techniques such as optical flow or binocular stereo. *Model-based* techniques are generally termed 'shape from shading' although we will focus on photometric stereo. We also consider alternative approaches which do not explicitly determine surface geometry. These are *appearance-based* methods which use mathematical techniques to provide compact representations of a set of images of the surface under various illumination directions. The relighting procedure depends on the specific technique but is generally achieved by manipulation of the resulting coefficients.

2.4.1 Criteria

Before considering either appearance-based or model-based relighting, we re-cap the criteria stated in Section 1.3 which will be used to identify the techniques most suited to our needs. The criteria are:

- 1. Suitable for globally flat diffuse surfaces.
- 2. Use of consumer-level equipment.
- 3. Practical input data capture.
- 4. Computationally efficient generation of representation.
- 5. Compact and compatible representation.
- 6. Computationally efficient generation of accurate relit images.

The ideal technique uses inexpensive equipment, requires few input images and minimal calibration. The resulting representation is fast to compute, compact, compatible with programmable graphics cards and provides accurate relit images of diffuse surface textures.

2.4.2 Appearance-based Relighting

Shashua showed that a linear combination of three images of a Lambertian surface illuminated from different directions is sufficient to form arbitrarily lit images of the surface [Shashua1992]. This is equivalent to a 3-dimensional basis in which the basis vectors are actually images. Principal component analysis and singular value decomposition have been used extensively to find such low-dimensional representations for large sets of images of a surface or object [Turk1991, Hallinan1994, Epstein1995, Belhumeur1998, Georghiades1999, Lin1999, Nishino2001, Ramamoorthi2002, Dong2003a]. Depending on the specific application (e.g. recognition), the images sample a range of illumination or viewing directions. With regard to the former,

Epstein *et al* found that the first three *eigenimages* correspond to the diffuse component, the fourth to the specular lobe and the remainder to the specular spike, shadows and occlusions. Linear combinations of such basis images generate new images under arbitrary illumination.

Other mathematical techniques have been successfully utilised to express a set of sample images as linear and non-linear combinations of functions e.g. 9-dimensional spherical harmonics [Basri2001]. With specific reference to texture, Malzbender proposed the 6-dimensional representation of polynomial texture maps (PTM) which enable shadows and interreflections to be modelled [Malzbender2001]. However, PTMs are not suitable for modelling specular reflectance. Although more often concerned with viewpoint variation, image-based rendering has also been employed to relight surfaces with more general reflection properties but entails the acquisition of densely sampled image sets [Koudelka2001, Debevec2000].

The large size of the datasets resulting from measurement of the BRDF or the BTF means that it is impractical to use the raw data directly. Various representations have been proposed to compress the data. Basis functions such as spherical harmonics [Cabral1987][Westin1992], Zernike polynomials [Koenderink1996a], wavelets [Lalonde1997] and cosine lobes [Lafortune1997] have been used to *approximate* the BRDF but can require large numbers of coefficients especially for materials exhibiting specular reflectance. 3D textons were introduced to represent the BTF [Leung2001] but the dimensionality of the appearance vectors is extremely high and the computational cost of reconstruction is expensive [Tong2002]. With regard to the densely sampled datasets, singular value decomposition was used to generate a basis for the BTF; this required more than 150 eigenvectors [Koudelka2003].

2.4.3 Model-based Relighting

These methods require relatively fewer images than the appearance-based techniques previously described. The disadvantage is that the reflectance model approximates the actual surface reflectance. Shape from shading (SFS) algorithms attempt to utilise a *single* image for this task [Horn1986]. However, it is impossible to unambiguously infer surface geometry, which has two degrees of freedom, from one pixel intensity without restricting the solution in some way or making assumptions. It is noted that the SFS problem is compounded for a Lambertian surface of varying albedo. This implies solving Equation 2.2 for *three* unknowns if it is sufficient to determine the product of albedo and intensity as a single parameter. A number of approaches to the SFS problem

have been detailed in the literature and a good survey is available although poor results are reported for all of the algorithms [Zhang1999].

The use of additional images makes the problem tractable. Lambertian photometric stereo uses multiple input images to determine the two surface gradients and albedo [Woodham1980]. Unlike SFS, photometric stereo is not under-determined and is straightforward to solve in this case. With regard to acquiring input data, the object or surface is illuminated from several different directions and images are captured by a static camera. Each image corresponds to one light source direction. The fact that both camera and object remain static means that registration problems are avoided since the images are automatically aligned with each other.

Lambertian Photometric Stereo

Photometric stereo largely entails using at least three images and in some cases many more to cope with non-Lambertian conditions. Two-image Lambertian techniques have been proposed in the literature but necessitate assumptions such as constant albedo, smooth surfaces and linearisation of the reflectance map by a Taylor's series expansion thus limiting their application [Onn1990, Lee1993, Hansson2000, Gullón2002]. Legendre polynomials have been used to represent Lambertian surfaces and are used in an iterative version of photometric stereo which uses either two or three images as input [Kim1997]. McGunnigle presented a simplified scheme for three-image Lambertian photometric stereo which is valid under a specific light source distribution [McGunnigle1998]. Spence and Chantler determined the optimal lighting distribution for the 3-image case [Spence2003a].

When specular highlights or shadows are present in the input images Lambertian photometric stereo will inevitably produce less accurate results. One strategy is to make the method *robust* by identifying and eliminating or attenuating outlying pixel intensity values. Coleman and Jain use a fourth image which allows an albedo estimate for each of the four permutations of three lights; the lowest albedo value is taken to correspond to diffuse reflection and thus specular highlights are avoided [Coleman1982]. This four-image method was extended to detect both specular highlights and shadows for colour images [Petrou2001]. Rushmeier *et al* proposed using five images: at each pixel the extreme intensity values are discarded to avoid highlights and shadows [Rushmeier1997]. Schlüns and Wittig used the dichromatric reflection model to identify and discard the specular component based on its spectral profile [Schlüns1993]. Polarisers have also been utilised to attenuate highlights in a method to determine the

topography of paper [Hansson2000]. Schlüns suggests a way to solve photometric stereo in the presence of self shadows but distinguishing these from cast shadows is difficult [Schlüns1997]. Woodham attempted to detect non-local phenomenon such as cast shadows and interreflections to make the photometric stereo algorithm more robust [Woodham1994]. Algorithms have been developed to cope with point light source illumination when the light is close to both the camera and the diffuse surface [Iwahori1990/1992/1994, Clark1992/1999]. These differential photometric stereo techniques utilise a *moving* point light source. Cho and Minamitani used an iterative technique to obtain surface geometry under point illumination after thresholding the data to suppress specular highlights [Cho1993].

In the case of noisy images, shape and albedo were iteratively recovered by formulating photometric stereo in the framework of the linear Kalman filter [Zhang1997]. Lee and Kuo use an explicit surface model and develop a photometric stereo algorithm which is less sensitive to noise because it is a global method rather than a local one [Lee1993]. A sensitivity analysis was carried out to determine the errors in surface normal orientation due to measurement errors in the input data [Jiang1991]. A technique to evaluate the noise robustness of photometric stereo has also been presented [Schlüns1997].

Reflectance models other than Lambert's law have been used in the photometric stereo algorithm. Tagare and deFigueiredo implemented a version of photometric stereo for a class of diffuse non-Lambertian surfaces by constructing corresponding reflectance maps from a mathematical expression [Tagare1991]. 'Photometric sampling' which entails using uniformly distributed extended light sources enables the recovery of the surface geometry and reflectance parameters of Nayar's hybrid model for both diffuse and specular reflection [Nayar1990]. Saito et al capture images under uniformly distributed light source at a constant illumination slant angle, discard the outlying data and fit a sine to the remaining Lambertian data [Saito1996]. Not only is it then possible to use Lambertian photometric stereo but the parameters of the Phong specular component can also be fitted to the separated intensity highlight. Ikeuchi implemented photometric stereo for specular reflecting surfaces to obtain the surface orientation by using distributed diffuse light sources and assuming perfect specular reflection [Ikeuchi1981]. Solomon and Ikeuchi used Coleman and Jain's four-image technique to extract specular pixels; this data was then fitted to a simplified version of the T-S model [Solomon1996]. Kay and Caelli used non-linear regression techniques with a simplified T-S model and simultaneously estimated surface reflectance and surface geometry
although this is expensive in terms of processing time [Kay1995]. Christensen and Shapiro describe a procedure for colour photometric stereo which can employ various reflection models such as T-S [Christensen1994].

Uncalibrated Photometric Stereo

The aforementioned photometric stereo techniques can all be classed as calibrated because the light source directions are known. Uncalibrated photometric stereo techniques in which they are unknown have also been extensively researched. In this case image intensities are the only data input to the photometric stereo algorithm. Woodham et al considered images of a Lambertian surface illuminated from different unknown directions and concluded that six images were sufficient to determine a solution with three remaining degrees of freedom [Woodham1991]. However, the original uncalibrated photometric stereo (UPS) algorithm is generally attributed to Hayakawa [Hayakawa1994]. His method involves constructing an intensity matrix each column of which corresponds to an image of the illuminated Lambertian surface under different light source directions and then factorising it into initial estimates of the surface orientation and light direction matrices. Whilst these estimates form a possible solution, it is not unique and an ambiguity exists. This is an example of the generic bilinear estimation-calibration problem [Koenderink1997]. Hayakawa proposes two ways to reduce the ambiguity to an orthogonal transformation. Other strategies for reducing or resolving this ambiguity have subsequently been proposed such as using the integrability constraint [Yuille1997, Belhumeur1999], using a calibration object [Yuille1999] and estimating light source directions [Spence2003b]. This uncalibrated photometric stereo method was extended to specular reflection; this is actually advantageous in terms of ambiguity reduction. Polarisation techniques were utilised to this end initially [Drbohlav2001]. Detecting specular highlights and assuming perfect specular reflection allowed estimates of the viewpoint to be obtained. The consistent viewpoint constraint was thus devised and used in conjunction with the integrability constraint to resolve the ambiguity [Drbohlav2002, Drbohlav2003]. This constraint was subsequently utilised with the T-S reflection model to allow its application to a wider range of materials [Georghiades2003a/b/c].

Recently another strategy for implementing uncalibrated photometric stereo was proposed [Hertzmann2003]. This method proposes the use of a calibration sphere such that intensities on the object of interest can be matched with those of the sphere. This orientation consistency means that the surface normal can be determined.

20

2.4.4 Criteria Compliance

The ideal relighting technique defined by the six criteria is inexpensive in terms of both equipment cost and data requirement and has a compact, compatible representation which provides accurate relit images of diffuse surface textures. Having reviewed both image-based and appearance-based techniques, it is apparent that in reality a suitable compromise between expense and accuracy will have to be found. The BRDF, BTF and some image-based rendering methods are accurate but prohibitively expensive due to the specialist equipment and dense sampling required. The eigenimage technique is an appearance-based approach which is more suited to our needs. It can be used to approximate more complex reflectance such as specular highlights and shadows although this would require more basis images. For diffuse reflectance the basis is three eigenimages and is therefore compact; the required interpolation for relighting could hence be carried out in hardware i.e. with a graphics card. However, bump maps [Blinn1978] are much more commonly utilised in the field of computer graphics. The implication is that it is preferable to explicitly determine the surface in terms of its gradients p and q. Furthermore, the bump map is often used in conjunction with a colour texture map [Woo1999] which corresponds to the albedo image. This data can be provided by the shape from shading model-based techniques. In particular we consider Lambertian three-image photometric stereo to be the most suitable candidate for diffuse reflecting surfaces. We disregard the single SFS and two image methods due to the requisite assumptions and accuracy concerns. Not only does this technique provide compact data in the most compatible form, the sampling required is also sparse. The disadvantage is that the surface texture reflectance is approximated by the Lambertian model but as previously noted, this provides reasonable results for a large range of surfaces. Uncalibrated Lambertian photometric stereo is another potential candidate. In this case the only data input is a minimum of six intensity images; significantly the corresponding illumination directions are not required.

2.5 Lambertian Photometric Stereo

In this section we present the mathematical framework for variants of the Lambertian photometric stereo method. To facilitate this, we first introduce the definitions required and then discuss the inherent assumptions and limitations of the algorithm.

2.5.1 Definitions

Co-ordinate System

As previously mentioned, the photometric stereo method requires that both camera and surface texture are static. Images are captured under different illumination directions. The corresponding equipment set-up is depicted in Figure 2.5. The co-ordinate frame is defined so that the optical axis corresponds to the *z*-axis. The camera is hence positioned on the *z*-axis with its line of sight along the axis such that its CCD array is parallel to the surface texture. The texture is globally flat and lies in the *x*-*y* plane. The direction of the light which is a single moveable source is defined by two angles. The slant angle σ is the angle between the illumination vector and the *z*-axis. The tilt angle τ is the angle between the *x*-axis and the projection of the illumination vector onto the *x*-*y* plane.



Figure 2.5 Equipment set-up for photometric stereo image capture.

Illumination Vector

The illumination vector is defined in terms of the slant and tilt angles as follows:

$$\mathbf{b} = \begin{bmatrix} \cos\tau\sin\sigma, & \sin\tau\sin\sigma, & \cos\sigma \end{bmatrix}^T$$
(2.5)

Scaled Illumination Vector

The scaled illumination vector is defined as follows:

$$\mathbf{l} = [l_x, l_y, l_z]^T = \lambda \mathbf{b}$$
(2.6)

The magnitude of this vector is the light source intensity λ :

$$\mathbf{l} = \lambda \tag{2.7}$$

Surface Normal

We consider the surface texture to be composed of components called facets each one of which corresponds to a unique image pixel. The orientation of each facet is given by the surface normal **n** which is perpendicular to the facet plane. It is often written in terms of the surface gradients p and q. If the texture is described by a height function z(x,y), the surface gradients are given by:

$$p = \frac{\partial z(x, y)}{\partial x} \qquad \qquad q = \frac{\partial z(x, y)}{\partial y} \qquad (2.8, 2.9)$$

The surface normal is determined by the cross-product of two non-parallel tangents to the facet plane which can be written in terms of the surface gradients [Horn1986].



Figure 2.6 Geometry of surface facet.

The tangent vectors are written as follows by considering small steps in both the x and y directions as illustrated in Figure 2.6.

$$\mathbf{r}_{x} = [\partial x, 0, p \partial x]^{T} = [1, 0, p]^{T}$$
(2.10)

$$\mathbf{r}_{\mathbf{y}} = \begin{bmatrix} 0, & \partial y, & q \partial y \end{bmatrix}^T = \begin{bmatrix} 0, & 1, & q \end{bmatrix}^T$$
(2.11)

The cross-product of these tangents gives a vector perpendicular to the facet plane :

$$\mathbf{r}_{x} \times \mathbf{r}_{y} = [-p, -q, 1]^{T}$$
(2.12)

Normalising to obtain a unit vector, the surface normal is hence written as follows:

$$\mathbf{n} = \frac{\left[-p, -q, 1\right]^{T}}{\sqrt{p^{2} + q^{2} + 1}}$$
(2.13)

Scaled Surface Normal

The scaled surface normal s is defined as the unit surface normal n scaled by the diffuse

albedo.
$$\mathbf{s} = [s_x, s_y, s_z]^T = \rho \mathbf{n}$$
 (2.14)

The magnitude of the this vector is the diffuse albedo:

$$|\mathbf{s}| = \rho \tag{2.15}$$

2.5.2 Assumptions & Limitations

Illumination

The unit illumination vector **b** defined by Equation 2.5, which is known if the method is calibrated, is taken to be constant over the entire surface texture. This means that the incident light rays are parallel to each other. The underlying assumption is that the surface is illuminated by a point light source at infinity. If the light source is too close for this assumption to hold then illumination vectors corresponding to each facet would be required for the per-pixel calculation. An inverse square law regarding the illumination radius would also have to be taken into account. This is usually ignored for the point source at infinity or at least incorporated into the albedo term since the radius can be taken as constant across the surface.

Reflection

The Lambertian photometric stereo model assumes that every facet is illuminated such that the corresponding pixel intensities are positive in every image. In theory shadows therefore do not feature in the analysis. For real world textures this is unlikely to be the case. The extent to which shadows are encountered depends on the convexity and roughness of the surface and the light source position. Self shadowing occurs when the angle between illumination vector and facet surface normal is greater than 90°. Cast shadowing occurs when a relatively high part of the surface occludes the incident light ray thus preventing it from reaching a relatively lower part. Cast shadowing is therefore a non-local process because the pixel intensity corresponding to one facet depends on other neighbouring facets. Another non-local process which is disregarded is interreflection. In this case we assume that surface facets do not act as secondary light sources. In reality the light reflected from neighbouring surface facets will increase the intensity of a facet to which it is directed. More significant increases would be observed for specular reflection but the surface is assumed to be diffuse in nature.

Imaging

Image irradiance is taken to be proportional to surface texture radiance. We assume that the surface texture is orthographically projected onto the camera sensor such that the viewing vector is constant over the surface i.e. for each facet. The camera sensor is also assumed to be linear. This means that the light which enters the camera lens should have a direct linear relationship to the image pixel value. The overall implication is that the image intensity is only a function of surface facet orientation and the albedo.

2.5.3 Mathematical Framework

Calibrated 3-Image

Assuming Lambertian reflectance, three images of a surface under different illumination conditions are sufficient to uniquely determine both the surface orientation and an albedo term [Woodham1980]. For each pixel position (x,y) the following can be written using Equation 2.2:

$$\mathbf{i}(x,y) = \begin{bmatrix} i_1(x,y) \\ i_2(x,y) \\ i_3(x,y) \end{bmatrix} = \lambda \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \mathbf{b}_3 \end{bmatrix} \rho \mathbf{n}(x,y)$$
(2.16)

Writing the three unit illumination vectors as a 3×3 matrix **B** and expressing the albedo and surface normal terms as the scaled surface normal **s**, this becomes:

$$\mathbf{i}(x, y) = \lambda \mathbf{Bs}(x, y) \tag{2.17}$$

where
$$\mathbf{B} = \begin{bmatrix} b_{1,x} & b_{1,y} & b_{1,z} \\ b_{2,x} & b_{2,y} & b_{2,z} \\ b_{3,x} & b_{3,y} & b_{3,z} \end{bmatrix}$$

Since the technique is calibrated, the corresponding tilt and slant angles are measured and the illumination matrix \mathbf{B} is known. The intensity at the pixel position is also known for each illumination condition from the three images of the texture. The system of equations is solved by inverting the illumination matrix \mathbf{B} and multiplying by the intensity vector \mathbf{i} .

$$\lambda \mathbf{s}(x, y) = \mathbf{B}^{-1} \mathbf{i}(x, y) \tag{2.18}$$

Assuming the unknown light source intensity to be constant for each of the three illumination vectors, we determine the surface gradients and an albedo term α as follows:

$$p = -\frac{s_x}{s_z} \tag{2.19}$$

$$q = -\frac{s_y}{s_z} \tag{2.20}$$

$$\alpha = \rho\lambda = \lambda\sqrt{s_x^2 + s_y^2 + s_z^2}$$
(2.21)

It is sufficient to combine the diffuse albedo and the light source intensity into this single parameter because their individual values are usually not required [Klette1999].

On a practical note, singular value decomposition (SVD) is one technique which is commonly utilised to perform the inversion operation. Decomposing the matrix \mathbf{B} into two orthogonal matrices and a diagonal matrix facilitates this. This is because inverting an orthogonal matrix merely involves transposing it whilst inverting a diagonal matrix entails the inversion of each diagonal element.

$$\mathbf{B} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^{\mathrm{T}} \tag{2.22}$$

$$\mathbf{B}^{-1} = \mathbf{V}\boldsymbol{\Sigma}^{-1}\mathbf{U}^{\mathrm{T}}$$
(2.23)

Simplified Calibrated Scheme

McGunnigle proposed a simplified Lambertian photometric stereo scheme by using a specific configuration of illumination directions [McGunnigle1998]. Its development involves expressing Equation 2.2 explicitly in terms of the surface gradients and the illumination angles using Equations 2.5 and 2.13:

$$i(x, y) = \rho \lambda \frac{(-p \cos \tau \sin \sigma - q \sin \tau \sin \sigma + \cos \sigma)}{\sqrt{p^2 + q^2 + 1}}$$
(2.24)

If the surface is illuminated from tilt angles of 0° , 90° and 180° at constant slant angle σ , the following expressions are obtained from Equation 2.24.

$$i_0(x, y) = \rho \lambda \frac{(-p \sin \sigma + \cos \sigma)}{\sqrt{p^2 + q^2 + 1}}$$
(2.25)

$$i_{90}(x,y) = \rho \lambda \frac{(-q \sin \sigma + \cos \sigma)}{\sqrt{p^2 + q^2 + 1}}$$
(2.26)

$$i_{180}(x,y) = \rho \lambda \frac{(p \sin \sigma + \cos \sigma)}{\sqrt{p^2 + q^2 + 1}}$$
(2.27)

If the first and third equations are added, a non-linear function of the surface derivatives will be obtained.

$$i_{NL}(x,y) = i_0(x,y) + i_{180}(x,y) = \frac{2\rho\lambda\cos\sigma}{\sqrt{p^2 + q^2 + 1}}$$
(2.28)

Two linear functions are then obtained by dividing the Equations 2.25 & 2.26 with Equation 2.28. These equations relate the surface gradients to image intensity and are independent of the albedo ρ and the light source intensity λ .

$$i_p(x,y) = \frac{i_0}{i_0 + i_{180}} = \frac{-p\tan\sigma + 1}{2}$$
(2.29)

$$\dot{i}_q(x,y) = \frac{\dot{i}_{90}}{\dot{i}_0 + \dot{i}_{180}} = \frac{-q \tan \sigma + 1}{2}$$
(2.30)

Once these equations have been re-arranged, expressions for the surface derivatives and albedo product can be obtained.

$$p = \frac{1 - 2i_p}{\tan \sigma} \tag{2.31}$$

$$q = \frac{1 - 2i_q}{\tan \sigma} \tag{2.32}$$

$$\alpha = \rho\lambda = \frac{i_0\sqrt{p^2 + q^2 + 1}}{-p\sin\sigma + \cos\sigma}$$
(2.33)

Over-constrained Calibrated Scheme

If more than three images of the surface texture illuminated from different light source directions are available then an over-constrained version of Lambertian photometric stereo can be implemented [Woodham1980]. In this case an image intensity matrix \mathbf{I} is formed each column of which corresponds to a single image illuminated from a unique direction. If there are *m* pixels in the image and *f* frames i.e. images with a different illumination direction then the matrix is formulated as follows:

$$\mathbf{I} = \begin{bmatrix} i_{11} & \dots & i_{1f} \\ \vdots & \dots & \vdots \\ i_{m1} & \dots & i_{mf} \end{bmatrix}$$
 where $m = \text{no. of pixels}$
 $f = \text{no. of frames}$ (2.34)

Lambert's law as given for a single pixel position by Equation 2.2 can be re-formulated in matrix terms as follows:

$$\mathbf{I} = \mathbf{SL} \tag{2.35}$$

In this case S is the scaled surface normal matrix which is $m \times 3$ in size. Each row of S is the transpose of an individual scaled surface normal.

$$\mathbf{S} = \begin{bmatrix} \mathbf{s}_1 \\ \mathbf{M} \\ \mathbf{s}_m \end{bmatrix} = \begin{bmatrix} s_{1x} & s_{1y} & s_{1z} \\ \mathbf{M} & \mathbf{M} & \mathbf{M} \\ s_{mx} & s_{my} & s_{mz} \end{bmatrix}$$
(2.36)

The matrix **L** is the scaled illumination vector matrix and is $3 \times f$ in size. Each column of **L** corresponds to an individual scaled illumination vector.

$$\mathbf{L} = \begin{bmatrix} \mathbf{l}_{1} & \Lambda & \mathbf{l}_{f} \end{bmatrix} = \begin{bmatrix} l_{1x} & \Lambda & l_{fx} \\ l_{1y} & \Lambda & l_{fy} \\ l_{1z} & \Lambda & l_{fz} \end{bmatrix}$$
(2.37)

Since f > 3 in this over-constrained case, **L** is not a square matrix. However, SVD can be employed to find the *pseudo*-inverse and hence provide a least-squares solution of the photometric stereo scheme for the scaled surface normal matrix.

$$\mathbf{S} = \mathbf{I}\mathbf{L}^{-1} \tag{2.38}$$

In reality it is the unit illumination matrix **B** which is known rather than scaled illumination matrix **L**. The estimate of the scaled surface normal matrix **S** will be scaled by the light intensity λ . The only implication is that the resulting albedo image corresponds to values of the albedo term α as previously discussed.

Uncalibrated Photometric Stereo

The mathematical framework for uncalibrated photometric stereo is analogous to the over-constrained calibrated technique and is given by Equation 2.35. Only the input intensity matrix \mathbf{I} is known, however. In this case the input intensity matrix is factorised into a *pseudo* surface matrix and a *pseudo* illumination matrix, $\hat{\mathbf{S}} \& \hat{\mathbf{L}}$. These represent a possible solution but it is not unique and an ambiguity exists since the following holds:

$$\mathbf{I} = \mathbf{\hat{S}}\mathbf{\hat{L}} = \mathbf{\hat{S}}\mathbf{A}^{\mathrm{T}}\mathbf{A}^{-\mathrm{T}}\mathbf{\hat{L}}$$
(2.39)

Even though the initial decomposition to obtain the first estimates is straightforward in terms of mathematics, the determination of the ambiguity matrix \mathbf{A} such that the resulting solutions will be unique is more challenging. This aspect of the uncalibrated technique will be covered in both Chapters 5 & 6. We note that it is common practice in the literature to use the *transpose* of \mathbf{A} in Equation 2.39 [Drbohlav2002, Yuille1999]. This notation proves to be convenient for the development of the relevant theory.

Uncalibrated Photometric Stereo with Outliers

If the input intensity data corresponds to a diffuse reflecting surface but contains outliers corresponding to shadows and highlights then the factorisation in the uncalibrated photometric stereo technique can be modified to attenuate the outlying data [Georghiades2003a]. This involves setting a lower and upper bounds on the intensity value; this can be based on the mean intensity plus or minus a specified number of standard deviations. The original intensity matrix **I** is analysed in order to generate two vectors, $\mathbf{v_r} \& \mathbf{v_c}$, which contain the indices of valid rows and valid columns i.e. those which have no intensity values outwith the set range. A reduced intensity matrix of completely valid rows is formed and decomposed with SVD in order to obtain an initial estimate of the pseudo illumination matrix:

$$\mathbf{I}^{Vr} = \hat{\mathbf{S}}^{Vr} \hat{\mathbf{L}}$$
(2.40)

where Vr denotes indices of rows with no invalid intensities. An estimate of the pseudo scaled surface normal matrix $\hat{\mathbf{S}}$ is then generated row-wise using $\hat{\mathbf{L}}$ and the original intensity matrix \mathbf{I} . This is achieved by taking the i^{th} row of \mathbf{I} and removing the invalid intensities given by the indices in \mathbf{v}_c to create a row vector \mathbf{i}_i^{Vc} . The illumination matrix is reduced accordingly and its pseudo-inverse determined such that the vector $\hat{\mathbf{s}}_i$ can be estimated:

$$\mathbf{i}_{i}^{Vc} = \hat{\mathbf{s}}_{i} \hat{\mathbf{L}}^{Vc} \xrightarrow{svd} \hat{\mathbf{s}}_{i} = \mathbf{i}_{i}^{Vc} \left(\hat{\mathbf{L}}^{Vc} \right)^{\dagger}$$
(2.41)

where *Vc* denotes indices of columns with no invalid intensities and the superscript \dagger denotes the pseudo-inverse. Next an updated estimate of the pseudo illumination matrix $\hat{\mathbf{L}}$ is generated column-wise using $\hat{\mathbf{S}}$ and the original intensity matrix. This is achieved by taking the *j*th column of **I** and removing the invalid intensities given by the indices in $\mathbf{v_r}$ to create a column vector \mathbf{i}_j^{Vr} . The scaled surface normal matrix is reduced accordingly and its pseudo-inverse determined such that the illumination vector \mathbf{l}_j can be estimated:

$$\mathbf{i}_{j}^{Vr} = \hat{\mathbf{S}}^{Vr} \hat{\mathbf{l}}_{j} \xrightarrow{svd} \hat{\mathbf{l}}_{j} = \left(\hat{\mathbf{S}}^{Vr} \right)^{\dagger} \mathbf{i}_{j}^{Vr}$$
(2.42)

This forms the basis for an iterative procedure. It is repeated until the estimates of both pseudo matrices converge.

2.5.4 Summary

Calibrated Lambertian photometric stereo can be solved with three input images by inverting the known illumination matrix. If the illumination distribution is restricted such that the slant angle is constant and the difference in tilt angle between successive directions is 90° then the scheme simplifies such that a combination of the input image intensities is merely required. If the problem is over-constrained, a least squares solution is found by finding the pseudo-inverse of the non-square illumination matrix. If the illumination direction corresponding to each image is unknown the uncalibrated technique may be employed although extra information is required to resolve the inherent ambiguity in the solution. If outlying data is present it is possible to use an iterative technique to attenuate it.

It is straightforward to utilise the resulting scaled surface normals, which may be formulated as p and q maps and albedo image, for relighting using the Lambertian model. For graphics applications it would also be possible to use the same data with the Phong model to provide hybrid reflectance if the user specifies the specular parameters.

2.6 Summary & Discussion

In this chapter we initially considered the processes which take place when light is reflected from a surface and the ways in which this can be modelled. We subsequently presented a survey on techniques to determine surface texture representations which facilitate relighting under arbitrary illumination conditions. We selected a group of methods based on defined criteria. Hence we identified Lambertian photometric stereo methods as the most appropriate with regard to fulfilling our relighting objective.

Photometric stereo methods are inexpensive with regard to equipment since an ordinary CCD camera can be utilised for image capture. Furthermore, they are inexpensive with regard to input data since extremely sparse sampling is merely required. The resulting representation for relighting which we required to be of low dimension is in terms of the two surface gradients and the albedo. Although the appearance-based eigenimage method also results in a 3-dimensional basis, relighting would involve linear combinations of the basis images. Whilst this procedure is certainly compatible with graphics hardware, bump mapping is much more commonly utilised in computer graphics applications. We have therefore opted for the photometric stereo class of methods because of the resulting surface-explicit representation. The mathematical framework required to implement several variants of Lambertian photometric stereo was presented.

Chapter 3

Surface Texture

3.1 Introduction

In the previous chapter we identified several methods which facilitate the relighting of surface textures under arbitrary lighting conditions. In order to assess the various techniques, intensity images of an illuminated surface are required as data input. Images of a sample surface texture can be captured with a CCD camera. In addition to utilising real world textures it is advantageous to have the ability to evaluate each technique under controlled conditions. Simulations can be performed with synthetic texture surfaces. This entails the use of mathematical models to generate data with user-specified surface texture characteristics. This approach allows favourable operating conditions to be determined and potential performance issues with real world textures to be predicted. In this chapter both synthetic and real surface textures are introduced. Their characteristics are described with a view to explaining the reason for their inclusion in the database of texture images used for the work presented in this thesis.

This chapter is organised as follows:

A working definition of surface texture is given in Section 3.3. A number of quantitative measures for surface description are given in Section 3.3. The synthetic and real world surface textures utilised as data input in this thesis are presented and characterised in Sections 3.4 and 3.5. A means of algorithm performance assessment which is based on a comparison of data set images and corresponding relit images is given in Section 3.6. We summarise the chapter in Section 3.7.

3.2 A Definition of Surface Texture

The term 'surface texture' has no precise definition and can be interpreted in a number of different ways depending on the application. It can refer to the two-dimensional intensity variation of the surface reflectance. This variation in albedo could be random or regular such that a basic pattern is repeated. Texture can also refer to the threedimensional structural features of a surface which may influence its feel as well as its appearance. In this case it is the topography or surface relief which is relevant. Smith et al note that surface texture therefore constitutes a 2D photometric property, a 3D geometric surface property or a combination of both [Smith2000]. The visual perception of such phenomena will differ depending on the scale at which the surface is observed. With regard to surface geometry, what appears smooth at one scale will be coarse at another. Koenderink et al defined three surface structure classes to facilitate a more informative description: they are the megastructure, mesostructure and microstructure [Koenderink1996b]. Megastructure refers to the object's global shape, mesostructure represents topographic texture i.e. local variation in surface height whilst microstructure refers to the features visible at relatively high magnification. We consider the term 'surface texture' or '3D surface texture' to denote a surface with mesostructural features whose megastructure is planar. Since the albedo may also be variable, our data sets hence comprise images of globally flat rough patterned surfaces (see Figure 3.1).



Figure 3.1 The two components of the surface texture definition (a) colour albedo and (b) surface relief.

3.3 Surface Description

A statistical approach is commonly adopted in order to describe the mesostructure of rough surface textures. In this case surface height is regarded as a two-dimensional random field. This means that standard statistics and signal processing techniques can be applied to characterise the texture.

Surface roughness

The most common measure of surface roughness is root mean square roughness z_{rms} . It is the standard deviation of height along a surface profile or over an area. The surface gradients p and q determined by application of photometric stereo can be integrated to provide an estimate of surface height in the form of a two-dimensional intensity image [Frankot1988, Gullón2002]. In this case surface roughness is given by the following equation for a square image:

$$z_{rms} = \sqrt{\frac{1}{m} \sum_{x=1}^{\sqrt{m}} \sum_{y=1}^{\sqrt{m}} [z(x, y) - \overline{z(x, y)}]^2}$$
(3.1)

where z(x, y) is the mean surface height, z(x, y) is the surface height at a point (x, y) in the image and *m* is the number of discrete height estimates which are equally spaced.

A less common approach is to use the root mean square slope. It is calculated in both directions as follows for square images of each gradient:

$$p_{rms} = \sqrt{\frac{1}{m} \sum_{x=1}^{\sqrt{m}} \sum_{y=1}^{\sqrt{m}} [p(x, y) - \overline{p(x, y)}]^2}$$
(3.2)

$$q_{rms} = \sqrt{\frac{1}{m} \sum_{x=1}^{\sqrt{m}} \sum_{y=1}^{\sqrt{m}} [q(x, y) - \overline{q(x, y)}]^2}$$
(3.3)

It is noted that both values of rms roughness and rms slope are affected by characteristics of the surface height image e.g. sampling frequency/resolution, sample length. This is particularly true for the rms slope which depends on both height amplitude and spacing.

Directionality

Despite this sensitivity, the rms slope can be used to provide an indication of the directionality of the surface texture [Gullón2002]. The directionality is a ratio defined as follows: $d = -\frac{p_{rms}}{rms}$

$$T = \frac{p_{rms}}{p_{rms} + q_{rms}}$$
(3.4)

If the texture is isotropic such that the surface is equally rough in all directions both rms slope values should have approximately the same value and d will tend to 0.5. If the texture is directional in either the x or y axis then d will tend towards 0 and 1 respectively.

Second Order Statistics

The single value parameters introduced above are unable to quantify the relationship between the values of the surface height field and are therefore of limited use. This is also true of the first order probability density function or histogram; if it can be described by a Gaussian this is simply equivalent to rms roughness. Second order statistics provide a measure of the correlation between pixels, however. The autocorrelation function (ACF) can be thought of as the overall intensity which results from overlaying two transparencies of the height image with one image shifted by a distance. Its Fourier transform equivalent is the power spectral density function (PSD) and facilitates frequency-based texture characterisation. For example, prominent peaks in the Fourier spectrum indicate the principal directionality of global texture patterns whilst the peak location gives the fundamental spatial period of the patterns. The PSD will be considered further in the following section.

3.4 Synthetic Surfaces

If the statistical representation of a surface texture or its spectral equivalent is sufficiently detailed it can be employed as a model. The three synthetic surfaces utilised in our simulation work are defined in terms of their PSD. It is noted that in order to describe an image completely, both the PSD and the phase spectrum are required. The phase spectrum contains information pertaining to edges i.e. boundaries of abrupt intensity change. We model the phase spectrum as an uncorrelated random field with a normal distribution. This means that the generated surfaces will be unstructured and as such they will have the appearance of natural rough surfaces.

3.4.1 Surface Models

The use of the following three surface models was inspired by research undertaken by Linnett and McGunnigle [Linnett1991, McGunnigle1998].

Fractal

Fractal Brownian functions provide a good model for describing rough natural surfaces [Pentland1988]. An important property of these surfaces is that they are statistically self-similar from one segment of the surface to another and at different scales. The fractal dimension D is related to the power roll-off factor β for the PSD of a fractal surface which is given as follows:

$$s_f(\omega) = \frac{k_f}{\omega^{\beta}} \tag{3.5}$$

where $\beta = 8-2D$, ω is the frequency and k_f is a constant. We used a value of 2.5 for the fractal dimension to produce a surface which resembles a natural surface. The power roll-off value β was 3.0. Surface roughness can be controlled by altering the value of k_f which is related to the height variance:

$$z_{rms} = \int_{-\infty}^{\infty} s(\omega) d\omega \tag{3.6}$$

Mulvaney

This PSD generates a rough surface which appears to have undergone a degree of physical processing. Mulvaney *et al* achieved this by employing a cut-off frequency [Mulvaney1989]. The PSD is given by:

$$s_m(\omega) = k_m \left(\frac{\omega^2}{\omega_c^2} + 1\right)^{-\frac{3}{2}}$$
(3.7)

where ω_c is the cut-off frequency in cycles per image and k_m is a constant. When the frequency ω is small this corresponds to a white noise spectrum since the PSD tends towards the specified constant k_m . For higher frequency the spectrum becomes equivalent to a fractal with a roll-off value β of 3.0.

Sand ripple

The fractal model was extended by maintaining the isotropic frequency decay of its PSD but shifting the peak from the centre and duplicating it to preserve symmetry [Linnett1991]. PSD plots are provided in Figure 3.2 for comparison. This directional fractal model results in a surface which resembles sand ripples on the seabed. It is given by the following equation:

$$s_{s}(u,v) = \frac{k_{s}}{\left(\left(u - u_{c}\right)^{2} + \left(v - v_{c}\right)^{2}\right)^{\frac{\beta}{2}}}$$
(3.8)

where *u* and *v* are the Cartesian frequency coordinates, u_c and v_c are cut-off frequencies in the *x* and *y* direction and k_s is a constant. The values of u_c , v_c and β were 64.0, 0.0 and 3.0 respectively.



Figure 3.2 PSD plot corresponding to (a) fractal & (b) sand ripple surfaces.

An inverse Fourier transform is performed to obtain the height maps for each synthetic surface (see Figure 3.3).



Figure 3.3 Height map images corresponding to (a) fractal, (b) Mulvaney, (c) sand ripple surfaces.

3.4.2 Data Generation

The height maps were used to generate images of each synthetic surface under userspecified illumination conditions and reflectance model. Sample images produced using the Lambertian model are given in Figure 3.4.



Figure 3.4 Images of the (a) fractal, (b) Mulvaney, (c) sand ripple surfaces illuminated from two directions (indicated by arrow) using the Lambertian model.

Every data set consists of 108 images which were generated with a set range of illumination directions for each synthetic surface using a selected reflectance model. These images correspond to a complete revolution with regard to illumination tilt angle with a $\Delta \tau$ of 10° at a constant slant angle for three slant angles of 30°, 45° and 60°. To expedite the processing of multiple data sets with several versions of the algorithms, an image size of 128 × 128 pixels was used in each case. An exception to this was when specific images were required for display in which case 512 × 512 pixels was used to provide sufficiently high resolution.

3.4.3 Characterisation

In this section we consider the character of the synthetic surface textures used to generate the data sets with a view to explaining the reason for their inclusion in this investigation. To do so, we examine reflectance characteristics and surface features corresponding to mesostructure and megastructure. Whilst the images have been produced under user-specified and hence known conditions in this case, we also establish various methods of describing texture character which could be employed on a general basis.

Synthetic Surface	\overline{p}	\overline{q}	Zrms	d
fractal	0.00	0.00	1.80	0.50
Mulvaney	0.00	0.00	0.62	0.50
sand ripple	0.00	0.00	0.44	0.84

Table 3.1 Measures of surface structure features for the three synthetic surfaces.

Reflectance

The majority of synthetic texture data sets were produced by relighting the height maps with the Lambertian model. A finite difference method was utilised to produce the pand q maps from the height map to facilitate this [Kreyszig1983, Chap.19]. Constant diffuse albedo was used in each case. This allows a visual inspection of the effect of any sources of error introduced since a recovered albedo image should ideally be of constant intensity. Negative intensities were not permitted; any encountered were set to zero in order to model self shadows. Data sets were also generated by using alternative reflectance models to provide deviation from Lambertian behaviour. Point lighting and Phong models which were discussed in the previous chapter were used for this purpose. Noise was also added to images produced by relighting with the Lambertian model to create noisy data sets.

Surface Mesostructure

In this case we used the rudimentary roughness and directionality measures described in Section 3.3 to provide a quantitative assessment of surface relief. The resulting values are given in Table 3.1. The fractal is the roughest of the three surfaces; both the Mulvaney and sand ripple textures are not rough surfaces according to this measure, however. Data sets of rougher surface textures were created by generating height maps for the three surfaces with increased height variance. With regard to directionality, the values of d indicate that the both the fractal and the Mulvaney surface are isotropic in nature. The value for the sand ripple surface means that it is a directional texture with the greatest deviation in surface gradient along the x-axis. Fourier spectra provide a

visual and more informative insight into the directionality of each of the synthetic surfaces compared to the value of d. Spectra corresponding to the illuminated surfaces of Figure 3.4 are given in Figure 3.5.



Figure 3.5 Fourier spectra of images of the (a) fractal, (b) Mulvaney, (c) sand ripple surfaces illuminated from two directions $\tau=0^\circ$, 90°. (indicated by arrow).

The spectra corresponding to images of the illuminated sand ripple surface confirm that it is a directional surface (Figure 3.5c) The peaks are characteristic of a surface with a single strong dominant frequency. The spacing of the peaks provides an indication of the period; the further apart they are the higher the frequency. The prominent peaks do not rotate with changing illumination direction although their intensity varies. This can be seen more clearly from the corresponding polar plot in Figure 3.6b. The polar plot is determined from the Fourier spectrum expressed in polar coordinates.

The spectra of both the fractal and the Mulvaney illuminated surfaces are typical of isotropic surfaces (Figure 3.5a/b). In contrast to directional surfaces, their spectra rotate with the illumination direction. This can also be seen from the polar plot in Figure 3.6a. Whilst the surface gradients are not aligned with any particular direction in this case, the *images* of illuminated isotropic surfaces are directional. This means that image directionality for isotropic surfaces depends on the illumination. Indeed this effect can be faintly detected upon observation of the images corresponding to the fractal and Mulvaney textures in Figure 3.4.



Figure 3.6 Polar plot of (a) an isotropic texture and (b) a directional texture illuminated from directions with different tilt angles.

The fact that the directionality of illuminated isotropic surfaces depends on the light source position suggests that given such an image, it would be possible to estimate illumination direction. This has been investigated by several authors [Knill1990, Chantler1997, Koenderink2003, Varma2004]. Chantler's technique operates in the frequency domain. He applies Fourier analysis to determine the tilt angle which corresponds to the peak value of the polar plot (see Figure 3.6a). A number of other methods for estimating the illumination direction in scenes have been presented but are not applicable in these circumstances due to the underlying assumptions e.g. [Pentland1982, Zheng1991, Chojnacki1994].

Surface Megastructure

The surface textures are assumed to have a planar megastructure i.e. globally flat. This can be checked by computing the mean surface gradient in both x and y axis directions. If the surface is flat then the value of both means should be zero; this implies that the mean surface normal should be $[0,0,1]^{T}$. According to the values given in Table 3.1 this is the case for all three synthetic surfaces.

3.5 Real World Textures

The synthetic texture surfaces were modelled with uniform albedo to facilitate a visual inspection of the effect of sources of error. Most textures encountered in practice are likely to have variable albedo, however. This might be due to the pattern of a textile but in reality even subtle changes in tone will be enough to cause variability in the albedo value. Since we are concerned with relighting applications in areas such as garment design with textiles, it is also paramount to consider not just variable albedo but colour. The selection of real textures for our experiments reflects the need to investigate these characteristics and the issues associated with them.

3.5.1 Data Generation

We captured images of sample surface textures using a charge coupled device (CCD) camera manufactured by Vosskühler. The model is a CCD-1300 and has a resolution equivalent to 1280×1024 square pixels. The camera is used in conjunction with a framegrabber and generates 12 bit images in various TIFF formats. These formats include separate images corresponding to red (R), green (G) and blue (B), a single RGB image or a greyscale image. Although the camera is based around a single CCD, it utilises a three colour mosaic filter such that each pixel position corresponds to either red, green or blue. The filter is arranged in a Bayer pattern [Bayer1976] whereby green is sampled at a higher rate than red and blue; this is because it is taken to correspond to luminance. The implication is that the formats which generate colour images have been interpolated. For image processing purposes it is preferable to use only directly measured intensities and hence each image was saved in the greyscale format.

With regard to equipment set-up, the camera is suspended approximately 0.6m above the texture sample which is approximately 0.2m square. Both camera and texture are fixed in position. The light source is a 7W compact fluorescent light bulb. The light is fixed to a moveable arm which is adjusted manually to provide the required tilt and slant angles; the light is approximately at a distance of 0.6m from the centre of the texture sample.

Every data set consists of 108 images which were captured with a set range of illumination directions for each sample texture. These images correspond to a complete revolution with regard to illumination tilt angle with a $\Delta \tau$ of 10° at a constant slant angle for three slant angles of 30°, 45° and 60°.

3.5.2 Variable Colour Albedo

Although the selected image format is greyscale, each intensity value corresponds to one of the three colours in the mosaic filter. Each greyscale image hence contains colour information. Since the red, green and blue pixels are arranged in a known way (the Bayer array in this case: see Figure 3.7), it is straightforward to generate a corresponding colour image by using demosaicking methods such as interpolation. However, it is both convenient and advantageous in terms of accuracy to retain the single greyscale format for image processing purposes. This not only avoids the need for multiple application of an algorithm to each colour channel but ensures that the processing is based solely on measured image intensities rather than a mixture of measured and estimated values. Hence unnecessary errors are not introduced by the demosaicking process. Such interpolation algorithms are employed only when display of a colour image is required.

The demosaicking process basically entails the generation of a high resolution colour image which is based on the greyscale original. At each pixel position in the original, the intensity value for one colour is known. In order to produce a high resolution colour image, an intensity value corresponding to each of the other two colours must be estimated at the same position over the entire image to give three complete colour planes. This can be carried out with a variety of interpolation methods, a comprehensive overview of which is given by Ramanath *et al* [Ramanath2002]. Bearing in mind that we only require colour images for display and that these are not used to determine performance figures, we employed a simple but effective algorithm to effect the conversion. Bilinear interpolation involves averaging the intensity values available for neighbouring pixels of the colour which is unknown at the pixel location being considered.

R ₁	G ₂	R ₃	G ₄
G ₅	B ₆	G ₇	B ₈
R ₉	G ₁₀	R ₁₁	G ₁₂
G ₁₃	B ₁₄	G ₁₅	B ₁₆

Figure 3.7 Illustration of Bayer array for the top left corner of a camera greyscale image format.

At a pixel of known red or blue intensity, the green value is estimated by taking the mean of the four green intensity values of adjacent pixels, e.g. $G_6 = (G_2 + G_5 + G_7 + G_{10})/4$, $G_{11} = (G_7 + G_{10} + G_{12} + G_{15})/4$. At a pixel of known green intensity, the red or blue value is estimated by taking the mean of the two adjacent red or blue values in the same row, e.g. $R_2 = (R_1 + R_3)/2$, $B_7 = (B_6 + B_8)/2$. At a pixel of known red or blue intensity, the blue or red intensity is estimated by taking the mean of the four blue or red intensity is estimated by taking the mean of the four blue or red intensity is estimated by taking the mean of the four blue or red intensity values of adjacent pixels, e.g. $R_6 = (R_1 + R_3 + R_9 + R_{11})/4$, $B_{11} = (B_6 + B_8 + B_{14} + B_{16})/4$.



Figure 3.8 Image planes obtained by bilinear interpolation which are smaller in height & width due to form of the calculations.

The resulting interpolated colour planes are each two pixels shorter in length for both row and column due to the form of the calculations (see Figure 3.8). A colour sample of each of the thirty-one surface textures utilised in this investigation is given in Figure 3.9. These have been generated from 512×512 pixels images cut from the original greyscale images.



Figure 3.9 One image from each of the real world texture data sets utilised in the investigation (illumination direction corresponds to $\tau=0^\circ$, $\sigma=45^\circ$).



Figure 3.9 One image from each of the real world texture data sets utilised in the investigation (illumination direction corresponds to $\tau=0^\circ, \sigma=45^\circ$). (continued)

3.5.3 Characterisation

The thirty-one surface textures selected consist of textiles (20), materials which could be used in an architectural 3D rendering of a building exterior or interior (7) and miscellaneous textures such as LegoTM bricks (4). In this section we examine their character in an analagous way to the approach used for the synthetic surface textures. In this case the three-image Lambertian photometric stereo technique was utilised in conjunction with an integration step [Gullón2002] to estimate the requisite data to calculate the various measures (see Table 3.2).

Texture	\overline{p}	\overline{q}	Z _{rms}	d
а	0.06	-0.02	8.75	0.61
b	0.07	-0.04	8.46	0.48
С	0.03	-0.03	6.90	0.47
d	0.07	-0.02	7.22	0.47
е	0.04	-0.02	7.26	0.50
f	0.03	0.03	8.64	0.52
g	0.07	-0.05	8.77	0.48
h	0.11	0.07	11.20	0.48
i	0.05	-0.02	6.24	0.55
j	0.06	-0.02	7.01	0.56
k	0.03	-0.01	8.56	0.46
l	0.12	-0.06	9.81	0.50
т	0.05	-0.02	8.14	0.50
п	0.07	-0.03	7.77	0.58
0	0.07	-0.01	10.97	0.47
р	0.12	0.02	13.93	0.41
q	0.02	-0.03	8.13	0.52
r	0.08	-0.02	6.94	0.51
S	0.09	-0.03	8.52	0.46
t	0.07	-0.03	8.32	0.47
и	0.06	-0.05	5.70	0.47
v	0.16	-0.05	7.30	0.50
w	-0.04	-0.04	8.09	0.73
x	0.15	-0.05	8.80	0.64
у	0.02	-0.03	8.94	0.54
Z.	0.04	0.02	9.95	0.47
aa	0.07	0.00	10.27	0.17
ab	-0.12	-0.08	15.53	0.66
ac	0.20	0.02	11.93	0.18
ad	-0.05	-0.04	9.68	0.47
ae	0.00	0.00	9.49	0.48

Table 3.2 Measures of surface structure features for the real world surface textures. (Bold entries are referred to in the text.)

Reflectance

The majority of the real world texture sample materials are assumed to exhibit diffuse reflection such that the Lambertian model can be utilised to produce relit images for comparison. Unlike the synthetic surface textures, many of these have a variable colour albedo and this must be taken into account when generating the relit images. Deviations from ideal Lambertian reflection are likely to be present or even prevalent in each image. The fact that noise is an inherent part of the image capture process, that the light source is not at infinity, that interreflections do take place and that shadowing will be a feature of such rough surfaces means that the reflectance model is an approximation. Four specularly reflecting textures (h, r, ac & ad) were also deliberately included in the data sets to provide examples of outlying behaviour. The extent to which each deviation occurs will determine how good an approximation the Lambertian model is.

Surface Mesostructure

The rms roughness measure is indicative of the extent to which each of the surface textures are likely to be prone to shadowing. The z_{rms} values given in Table 3.2 show that all of the selected real world textures are rough. Textures with relatively high values such as ab, p, ac, h & o are likely to deviate significantly from ideal Lambertian behaviour.

With regard to directionality, the selected textures are either isotropic or directional in nature. The measure d allows the directional textures which have a *single* dominant direction to be identified. According to the values given in Table 3.2 the textures aa & ac are therefore *uni-directional* with the major variation in surface gradient aligned with the y-axis whilst w, ab and x are uni-directional with regard to the x-axis. The remaining textures are either isotropic or multi-directional. The measure d cannot be used to distinguish between the two, however. Many of the textures are woven textiles where the warp and weft threads are visible and will have two dominant directions resulting in a value of d similar to that for isotropic textures. Directionality can be more precisely predicted by examining the corresponding Fourier spectra of the illuminated surface texture image. Examples of these are given in Figure 3.10. The characteristic of the isotropic surface texture i.e. spectra rotating according to illumination direction, reveals that textures u, ad & m are isotropic. The other textures are directional. The uni-directional surfaces result in spectra similar to that presented in Figure 3.10b. The synthetic sand ripple texture is another example of this. Examples of the spectra of multi-directional textures are given in Figure 3.10c.



Figure 3.10 Example images of the Fourier spectra of real world textures which exhibit (*a*) *isotropic*, (*b*) *uni-directional and* (*c*) *multi-directional behaviour. These correspond to textures u, w and f respectively.*

Planar Megastructure

Whilst the sample textures have all been arranged such they are globally flat, it is recognised that this may not quite be the case for a small sample of their overall image. As previously mentioned, if a texture is globally flat then the mean of the surface derivatives should be zero. The mean values of both *p* and *q* output images which are 512×512 pixels in size are reported in Table 3.2. They demonstrate some deviation from the ideal mean surface normal of $[0,0,1]^{T}$. This is particularly true for the textures ac, v, x, ab & h.

3.6 Relighting Assessment

Subsets of both the generated synthetic data sets and the captured data sets of real world textures are used as data input for the various photometric stereo algorithms used in this research. The output from each algorithm application can be used in conjunction with an illumination model to generate relit images of the surface texture under arbitrary illumination conditions. In order to assess the algorithm performance in each case, we generated relit images with the Lambertian model. We used a signal to relight error ratio *SER* as a measure of the difference in intensity values between an original image and a corresponding relit image. This is based on the signal to residue ratio metric [McGunnigle1998, Gullón2002] and is calculated as follows:

$$SER(j) = 10 \log \left[\frac{\operatorname{var}[\mathbf{I}(\tau_j, \sigma_j)]}{\operatorname{var}[\mathbf{I}(\tau_j, \sigma_j) - \mathbf{I}_{relight}(\tau_j, \sigma_j)]} \right]$$
(3.9)

where $I(\tau_j, \sigma_j)$ is a data set image, $I_{relight}(\tau_j, \sigma_j)$ is the corresponding generated image relit under the j^{th} illumination direction defined by the tilt angle τ_j and slant angle σ_j and var[I] is the variance of the intensity image.

We obtained an estimate of the overall relight accuracy for each texture by calculating *SER* values for a series of relit images and taking their mean. The *texture signal to relight error ratio TSER* is hence given by:

$$TSER_{tx} = \frac{1}{n} \sum_{j=1}^{n} [SER_{tx}(j)]$$
(3.10)

for a texture tx where n is the number of original-relit image pairs considered. In order to avoid bias, we note that the relit images used to calculate the *TSER* should be estimates of images in the database which have not been used as input data. In our experiments we frequently used input data corresponding to a common slant angle and calculated this metric using database images corresponding to the other two slant angles. In this case the *TSER* value is calculated from images corresponding to seventy-two illumination directions.

A mean *TSER* was also calculated by averaging the individual values for each texture utilised in the experiments. For the real textures this was calcuated as follows:

$$\overline{TSER} = \frac{1}{31} \sum_{tx=a}^{ae} [TSER_{tx}]$$
(3.11)

A visual inspection of the relight error was also provided through the use of *difference images*. These are generated by subtracting the relit image from its corresponding data set image and taking the magnitude of the result i.e. $| I(\tau,\sigma)-I_{relight}(\tau,\sigma) |$. This implies that relatively accurate regions will be dark and relatively inaccurate regions will be light.

3.7 Summary & Discussion

In this chapter we explicitly defined the term 'surface texture' as referring to globally flat rough patterned surfaces. We discussed various approaches used to describe such textures such as roughness and directionality measures. We introduced three synthetic surface models based on second order statistics which were employed to generate images under a set range of illumination conditions using a number of different illumination models. The synthetic texture data sets thus generated provide a means of assessing the various photometric stereo algorithms under a series of controlled conditions which are both favourable and adverse. We introduced the real world surface texture data sets which consist of the images of coloured texture samples illuminated under identical illumination directions to the synthetic data sets and described the equipment used in their capture. We discussed the issues involved in the processing/display of the resulting coloured images and considered the character of the corresponding surface textures. These surface texture data sets provide a means of assessing the various photometric stereo algorithms under real conditions. Finally the texture signal to relight error ratio (TSER) was introduced as means of determining a single measure of performance to assess the algorithms in a consistent quantitative way.

Chapter 4

Optimal Illumination for Three-Image Photometric Stereo

4.1 Introduction

In this chapter we focus on the three-image Lambertian photometric stereo technique which was introduced in Chapter 2. In this case not only is intensity data available but the corresponding illumination conditions under which each image was captured are also known. This implementation of photometric stereo is hence classed as a *calibrated* technique because the direction of the illumination is known.

Illumination direction has a significant bearing on the accuracy of photometric stereo. Woodham advocates maximising the illumination slant angle for optimal performance [Woodham1980] although he notes that its value is restricted in practice. This is due to the need to minimise the presence of shadows which are detrimental to performance. With regard to the relative position of the three light sources Woodham points out that a co-planar illumination arrangement should be avoided [Woodham1980]. However, the illumination tilt angles which correspond to optimal performance have not been reported in the literature. We devised a means of determining the optimal illumination configuration [Spence2003a]. We present its theoretical development, which is based on noise minimisation, in this chapter. An equivalent process based on empirical data is also given.

This chapter is organised as follows:

The three-image photometric stereo equations originally introduced in Chapter 2 are re-stated in Section 4.2 and current guidelines regarding optimal performance are reviewed. In Section 4.3 a sensitivity analysis is performed with a view to deriving overall noise expressions. An equivalent empirical approach is developed in Section 4.4. The noise expressions are minimised in Section 4.5 and the optimal operating conditions are reported. A practical assessment of the proposed optimal illumination configuration is given in Section 4.6. Conclusions are drawn in Section 4.7.

4.2 Three-Image Photometric Stereo

As briefly discussed in Chapter 2, Woodham demonstrated that three images of a surface under different illumination conditions are sufficient to uniquely determine both the surface orientation and an albedo term.



Figure 4.1 Equipment arrangement for three-image Lambertian photometric stereo.

A system of equations based on the Lambertian model is solved by inverting the illumination unit vector matrix. We re-state Equations 2.17 and 2.18 to illustrate this:

$$\mathbf{i}(x, y) = \lambda \mathbf{B} \mathbf{s}(x, y) \rightarrow \lambda \mathbf{s}(x, y) = \mathbf{B}^{-1} \mathbf{i}(x, y)$$

The surface gradients p & q and the albedo parameter α can then be determined (Equations 2.19-2.21):

$$p(x, y) = -\frac{s_x(x, y)}{s_z(x, y)} \quad q(x, y) = -\frac{s_y(x, y)}{s_z(x, y)}$$
$$\alpha(x, y) = \lambda \sqrt{s_x(x, y)^2 + s_y(x, y)^2 + s_z(x, y)^2}$$

Accuracy is an issue which was considered by Woodham in some depth [Woodham1980]. Reflectance maps were used to illustrate his main argument. He recommends dense iso-intensity contours for maximum accuracy. In this case a large change in intensity is attained for a small change in the surface gradient values, p & q. In other words it is desirable to maximise $\delta i/\delta p$ and $\delta i/\delta q$. Dense iso-intensity contours are achieved by increasing the value of the slant angle σ as demonstrated in Figure 4.2. In practice the slant angle is limited due to the adverse effect of the increasing presence of shadows.



Figure 4.2 Reflectance maps for a Lambertian surface with illumination tilt angle of 0° and varying slant angle.

Apart from maximising the slant angle, recommendations for the relative position of the three light sources with regard to the tilt angle are not apparent in the literature. This issue is referred to indirectly by Woodham when he points out that the scheme cannot be solved when the illumination vectors are arranged in a co-planar configuration [Woodham1980]. The resulting illumination matrix will be uninvertible in this case. For their two-image photometric stereo algorithm Lee and Kuo argue that the gradient direction of the reflectance map for one of the images should correspond to the tangential directions of the reflectance map of the other image [Lee1993]. They propose to achieve this by employing a difference of 90° between the illumination tilt angles. Gullón shows that the accuracy of her two-image techniques is more sensitive to tilt angle difference than the illumination arrangement position relative to a unidirectional surface and confirms that $\Delta \tau = 90^{\circ}$ is optimal [Gullón2002]. With regard to using more than two lights with linear photometric stereo Gullón argues that an even arrangement is optimal since it maximises the linear term. The fact that side lighting acts as a directional filter of the surface height function suggests that the signal to noise ratio could be maximised by distributing the illumination tilt angles equally through 360° [Chantler1995]. However, this has never been formally investigated with threeimage photometric stereo.

4.3 Noise Expression Derivation

In the previous section we argued that the recommended operating conditions for three-image photometric stereo are not sufficiently precise to guarantee accurate and reliable results. It is apparent that more detail regarding the optimal placement of the lights is required [Spence2003a]. In this section we develop the theoretical approach which was undertaken to investigate this issue. This involves deriving expressions for noise in the estimates of the scaled surface normals and using them to provide an overall figure or merit.

Sensitivity Analysis

Sensitivity analysis is a common approach used to gain an insight into the behaviour of a mathematical model such as photometric stereo and forms an important foundation for the work presented in this chapter. It is the study of how the variation in the output of a model can be apportioned to different sources of variation [Saltelli2000, Chap.1]. In the case of photometric stereo, the output is the estimate of the surface orientation in the form of the scaled surface normal. We propose that ascertaining its response to variation in the input, namely the intensity images and their corresponding illumination conditions, would be useful in achieving our objective of determining optimal operating conditions. Jiang and Bunke carried out a sensitivity analysis to examine the effect of measurement errors in the input data of photometric stereo [Jiang1991]. However, they did not consider the corresponding optimal illumination configuration. Furthermore, we employ a different approach to effect the sensitivity analysis [Spence2003a].

With regard to practical implementation, sensitivity analysis often takes the form of a sampling-based procedure during which the model is executed repeatedly over an extensive range of input conditions. We used this approach in order to produce empirical results and it will be discussed later in the chapter. For a purely theoretical treatment, however, we derive expressions for the sensitivity of each scaled surface normal element s_x , s_y , s_z with respect to changes in the input image intensities i_1 , i_2 , i_3 .

When the illumination vectors are not constrained to be of common slant angle, the illumination matrix formed from them depends on six parameters. With tilt angles τ_i and slant angles σ_i where *i*=1, 2, 3, the unit illumination vector matrix **B** is:

$$\mathbf{B} = \begin{bmatrix} \cos\tau_1 \sin\sigma_1 & \sin\tau_1 \sin\sigma_1 & \cos\sigma_1 \\ \cos\tau_2 \sin\sigma_2 & \sin\tau_2 \sin\sigma_2 & \cos\sigma_2 \\ \cos\tau_3 \sin\sigma_3 & \sin\tau_3 \sin\sigma_3 & \cos\sigma_3 \end{bmatrix}$$
(4.1)

The inverse of **B** is determined:

$$\mathbf{B}^{-1} = \frac{1}{\text{determinant}(\mathbf{B})} \text{adjoint}(\mathbf{B})$$
(4.2)

Substituting the inverse into Equation 2.18 provides expressions for each component of the scaled surface normals.

$$s_{x} = -\frac{\begin{pmatrix} (\sin\tau_{3}\cos\sigma_{2}\sin\sigma_{3} - \sin\tau_{2}\sin\sigma_{2}\cos\sigma_{3})i_{1} \\ + (\sin\tau_{1}\sin\sigma_{1}\cos\sigma_{3} - \sin\tau_{3}\cos\sigma_{1}\sin\sigma_{3})i_{2} \\ + (\sin\tau_{2}\cos\sigma_{1}\sin\sigma_{2} - \sin\tau_{1}\sin\sigma_{1}\cos\sigma_{2})i_{3} \end{pmatrix}}{k_{ov}}$$
(4.3)

$$s_{y} = -\frac{\begin{pmatrix} (\cos\tau_{2}\sin\sigma_{2}\cos\sigma_{3} - \cos\tau_{3}\cos\sigma_{2}\sin\sigma_{3})i_{1} \\ + (\cos\tau_{3}\cos\sigma_{1}\sin\sigma_{3} - \cos\tau_{1}\sin\sigma_{1}\cos\sigma_{3})i_{2} \\ + (\cos\tau_{1}\sin\sigma_{1}\cos\sigma_{2} - \cos\tau_{2}\cos\sigma_{1}\sin\sigma_{2})i_{3} \end{pmatrix}}{k_{ov}}$$
(4.4)

$$s_{z} = \frac{\begin{pmatrix} (\sin(\tau_{3} - \tau_{2})\sin\sigma_{2}\sin\sigma_{3})i_{1} \\ + (\sin(\tau_{1} - \tau_{3})\sin\sigma_{1}\sin\sigma_{3})i_{2} \\ + (\sin(\tau_{2} - \tau_{1})\sin\sigma_{1}\sin\sigma_{2})i_{3} \end{pmatrix}}{k_{ov}}$$
(4.5)

where

$$k_{ov} = \sin(\tau_3 - \tau_2)\cos\sigma_1\sin\sigma_2\sin\sigma_3 + \sin(\tau_1 - \tau_3)\sin\sigma_1\cos\sigma_2\sin\sigma_3 + \sin(\tau_2 - \tau_1)\sin\sigma_1\sin\sigma_2\cos\sigma_3$$
Differentiating Equations 4.3, 4.4 and 4.5 with respect to each of the three image intensities gives nine sensitivity expressions. These describe how sensitive the error in the estimated components of the surface normal (compared to the true surface normal) is to error in the intensity measurements.

$$\frac{\partial s_x}{\partial i_1} = -\left(\frac{\sin\tau_3\cos\sigma_2\sin\sigma_3 - \sin\tau_2\sin\sigma_2\cos\sigma_3}{k_{ov}}\right)$$
(4.6)

$$\frac{\partial s_{y}}{\partial i_{1}} = -\left(\frac{\cos\tau_{2}\sin\sigma_{2}\cos\sigma_{3} - \cos\tau_{3}\cos\sigma_{2}\sin\sigma_{3}}{k_{ov}}\right)$$
(4.7)

$$\frac{\partial s_z}{\partial i_1} = \left(\frac{\sin(\tau_3 - \tau_2)\sin\sigma_2\sin\sigma_3}{k_{ov}}\right)$$
(4.8)

$$\frac{\partial s_x}{\partial i_2} = -\left(\frac{\sin\tau_1\sin\sigma_1\cos\sigma_3 - \sin\tau_3\cos\sigma_1\sin\sigma_3}{k_{ov}}\right)$$
(4.9)

$$\frac{\partial s_{y}}{\partial i_{2}} = -\left(\frac{\cos\tau_{3}\cos\sigma_{1}\sin\sigma_{3} - \cos\tau_{1}\sin\sigma_{1}\cos\sigma_{3}}{k_{ov}}\right)$$
(4.10)

$$\frac{\partial s_z}{\partial i_2} = \left(\frac{\sin(\tau_1 - \tau_3)\sin\sigma_1\sin\sigma_3}{k_{ov}}\right)$$
(4.11)

$$\frac{\partial s_x}{\partial i_3} = -\left(\frac{\sin\tau_2\cos\sigma_1\sin\sigma_2 - \sin\tau_1\sin\sigma_1\cos\sigma_2}{k_{ov}}\right)$$
(4.12)

$$\frac{\partial s_{y}}{\partial i_{3}} = -\left(\frac{\cos\tau_{1}\sin\sigma_{1}\cos\sigma_{2} - \cos\tau_{2}\cos\sigma_{1}\sin\sigma_{2}}{k_{ov}}\right)$$
(4.13)

$$\frac{\partial s_z}{\partial i_3} = \left(\frac{\sin(\tau_2 - \tau_1)\sin\sigma_1\sin\sigma_2}{k_{ov}}\right)$$
(4.14)

As Equation 2.18 is linear, the noise in the scaled surface normal s can be simply derived from the sensitivities given by these equations. We note that the effect of inaccuracies in the measurement of the illumination angles is not considered here. As mentioned, this was previously investigated and reported by Jiang *et al* [Jiang1991].

See Appendix B for the equivalent but simpler case when the slant angle is common to the three illumination directions.

Noise in the Scaled Surface Normal

The variance¹ of a parameter x which is a function of two variables u and v can be shown to take the following general form (see Appendix C for proof):

$$\psi_x^2 = \psi_u^2 \left(\frac{\partial x}{\partial u}\right)^2 + \psi_v^2 \left(\frac{\partial x}{\partial v}\right)^2 + 2\psi_{uv}^2 \frac{\partial x}{\partial u} \frac{\partial x}{\partial v}$$
(4.15)

If the variables are independent then Equation 4.15 simplifies:

$$\psi_x^2 = \psi_u^2 \left(\frac{\partial x}{\partial u}\right)^2 + \psi_v^2 \left(\frac{\partial x}{\partial v}\right)^2$$
(4.16)

Equations 4.3, 4.4 and 4.5 demonstrate that each element of the scaled surface normal is a function of three variables i.e. the input intensity images. If we assume that the noise in each image is Gaussian and independent then a variation of Equation 4.16 can be used to predict the noise in each element of **s**. For example, the variance of the noise in s_x is given by the following:

$$\psi_{s_x} = \sqrt{\psi_{i_1}^2 \left(\frac{\partial s_x}{\partial i_1}\right)^2 + \psi_{i_2}^2 \left(\frac{\partial s_x}{\partial i_2}\right)^2 + \psi_{i_3}^2 \left(\frac{\partial s_x}{\partial i_3}\right)^2}$$
(4.17)

Despite being independent if the magnitude of the noise in each input image has approximately the same variance $\overline{\psi}_i$ then:

$$\psi_{s_x} = \overline{\psi}_i \sqrt{\left(\frac{\partial s_x}{\partial i_1}\right)^2 + \left(\frac{\partial s_x}{\partial i_2}\right)^2 + \left(\frac{\partial s_x}{\partial i_3}\right)^2} \tag{4.18}$$

In order to allow a completely theoretical analysis the formulas were re-arranged to make them independent of input noise. A noise *ratio* is now predicted for each of the scaled surface normal elements. These expressions describe the error in the scaled surface normal relative to the average error in the input intensity measurements.

$$\frac{\psi_{s_x}}{\overline{\psi}_i} = \sqrt{\left(\frac{\partial s_x}{\partial i_1}\right)^2 + \left(\frac{\partial s_x}{\partial i_2}\right)^2 + \left(\frac{\partial s_x}{\partial i_3}\right)^2} \tag{4.19}$$

$$\frac{\psi_{s_y}}{\overline{\psi_i}} = \sqrt{\left(\frac{\partial s_y}{\partial i_1}\right)^2 + \left(\frac{\partial s_y}{\partial i_2}\right)^2 + \left(\frac{\partial s_y}{\partial i_3}\right)^2} \tag{4.20}$$

$$\frac{\psi_{s_z}}{\overline{\psi_i}} = \sqrt{\left(\frac{\partial s_z}{\partial i_1}\right)^2 + \left(\frac{\partial s_z}{\partial i_2}\right)^2 + \left(\frac{\partial s_z}{\partial i_3}\right)^2} \tag{4.21}$$

Substituting Equations 4.6 - 4.14 into 4.19, 4.20 and 4.21 gives the full equation for the noise ratio of each element in **s**.

¹ The symbol ψ^2 is used here for variance instead of σ^2 to distinguish from illumination slant angle σ .

Single Figure of Merit

Given that noise is present in the input intensity images, it is apparent from the noise ratio expressions (see Equations 4.19 - 4.21 & 4.6 - 4.14) that the resulting level of noise in the scaled surface normal estimates depends on the illumination configuration. Our objective is to establish operating conditions which minimise the noise in the output in order to determine accurate estimates of the surface normal. The *optimal* illumination configuration can therefore be found by minimising each of the three noise ratios. However, a single objective function is required in order to implement an optimisation procedure. It is possible to formulate such a metric by taking into account the intended use of the output data. We have chosen to consider image-based rendering applications. The intensity of a relit pixel under arbitrary illumination can be expressed as follows using Lambert's law.

$$i_{relieht}(x, y) = s_x(x, y)\cos\tau\sin\sigma + s_y(x, y)\sin\tau\sin\sigma + s_z(x, y)\cos\sigma$$
(4.22)

Since the tilt and slant angles are specified in order to generate a relit image, the trigonometric terms in Equation 4.22 reduce to scalars. Hence the relit intensity is simply a weighted sum of the elements of s.

$$i_{relight}(x,y) = k_x s_x(x,y) + k_y s_y(x,y) + k_z s_z(x,y)$$
(4.23)

We therefore choose our figure of merit to be the variance of the sum of the s_x , s_y and s_z noise processes. We assume that these noise processes are highly correlated, each being a function of the three image noise processes. In this case the overall variance is simply given by the sum of the variances of the individual scaled surface normal elements. See Appendix C for standard statistics proof. Our single figure of merit is hence given by the following equation.

$$M_{rough} = \frac{\psi_{s_x} + \psi_{s_y} + \psi_{s_z}}{\overline{\psi}_i}$$
(4.24)

Substituting Equations 4.19, 4.20 and 4.21 into 4.24 this becomes:

$$M_{rough} = \sqrt{\left(\frac{\partial s_x}{\partial i_1}\right)^2 + \left(\frac{\partial s_x}{\partial i_2}\right)^2 + \left(\frac{\partial s_x}{\partial i_3}\right)^2 + \sqrt{\left(\frac{\partial s_y}{\partial i_1}\right)^2 + \left(\frac{\partial s_y}{\partial i_2}\right)^2 + \left(\frac{\partial s_y}{\partial i_3}\right)^2} + \sqrt{\left(\frac{\partial s_z}{\partial i_1}\right)^2 + \left(\frac{\partial s_z}{\partial i_2}\right)^2 + \left(\frac{\partial s_z}{\partial i_3}\right)^2}$$
(4.25)

It is then straightforward to substitute Equations 4.6 - 4.14 into 4.25 in order to provide the final expression. This is a function of illumination tilt and slant angles. Its form is ideally suited to an optimisation analysis. It can therefore be used to determine the optimal illumination configuration. It may also be utilised to gauge the potential performance of any given illumination arrangement.

Smooth Surface Simplification

Whilst our approach does not take shadowing into account and is effectively independent of the distribution of surface normals, we may make an additional simplification for smooth surfaces. If the surface slopes are low then following Kube and Pentland [Kube1988] we can use a MacLaurin's expansion² of Lambert's law and ignore the higher order terms.

$$i(x,y) \approx \rho \lambda(p(x,y)\cos\tau\sin\sigma + q(x,y)\sin\tau\sin\sigma - \cos\sigma) \left[1 - \frac{1}{2!}(p^2 + q^2) + \dots\right] \quad (4.26)$$

This is possible because when the slope angles are less than 15° then the surface gradients p, q << 1 and the term $(p^2+q^2) < 0.1$. In this case s_z tends to a constant i.e. the local albedo and can be ignored for the purposes of a sensitivity analysis.

$$s_z = \frac{\rho}{\sqrt{p^2 + q^2 + 1}} \approx \rho$$
 when $p,q << 1$ since $p^2 + q^2 < 0.1$ (4.27)

The figure of merit for a smooth surface is therefore:

$$M_{smooth} = \frac{\psi_{s_x} + \psi_{s_y}}{\overline{\psi}_i}$$
(4.28)

Substituting Equations 4.19 and 4.20 into 4.28 this becomes:

$$M_{smooth} = \sqrt{\left(\frac{\partial s_x}{\partial i_1}\right)^2 + \left(\frac{\partial s_x}{\partial i_2}\right)^2 + \left(\frac{\partial s_x}{\partial i_3}\right)^2} + \sqrt{\left(\frac{\partial s_y}{\partial i_1}\right)^2 + \left(\frac{\partial s_y}{\partial i_2}\right)^2 + \left(\frac{\partial s_y}{\partial i_3}\right)^2}$$
(4.29)

Substituting Equations 4.6, 4.7, 4.9, 4.10, 4.12 and 4.13 into 4.29 provides an alternative figure of merit which specifically applies to smooth surfaces. This will allow the optimal behaviour of both types of surfaces to be compared and contrasted.

4.3.1 Summary

In this section we derived noise expressions for three-image Lambertian photometric stereo. We used these expressions to formulate two figures of merit M_{rough} and M_{smooth} which are applicable to rough and smooth surface textures respectively. These will be

² Equivalent to the more general Taylor's series expansion with parameters set to zero

used in an optimisation procedure to determine the optimal illumination configuration (see Section 4.5).

4.4 Empirical Determination of Noise

It is difficult to calculate the absolute noise in the photometric stereo process. However, the *temporal* noise in the process may be easily estimated for both input and output data. This facilitates an empirical investigation equivalent to the theoretical treatment detailed in the previous section (see Figure 4.3).



Figure 4.3 Flowchart depicting empirical sensitivity analysis.

With regard to data, images of a real isotropic texture (texture m) were acquired over a range of 86 illumination directions. A set of ten images was captured for each illumination direction. This meant that we could apply the photometric stereo algorithm ten times for a given illumination configuration. Temporal noise estimates corresponding to this configuration can be determined from the multiple input and output images. This is achieved by estimating the value of statistics parameters i.e. mean and variance for each set of ten images in order to calculate the mean standard deviation.

Noise in the Input Intensity Images

Given ten images $i^t(x,y)$ where t=1-10 of the texture which correspond to a single illumination direction, a temporal noise value is estimated as follows:

$$\psi_{i} = \frac{1}{m} \sum_{x=1}^{\sqrt{m}} \sum_{y=1}^{\sqrt{m}} \left(\frac{1}{9} \sum_{t=1}^{10} \left(i^{t} - \bar{i}(x, y) \right)^{2} \right)$$
(4.30)

where *m* is the number of pixels in a square image.

Since there are three illumination directions in the input data, three temporal noise values are calculated $\psi_{i_1}, \psi_{i_{21}}, \psi_{i_3}$ using Equation 4.30. A mean value of the input noise is estimated:

$$\overline{\psi}_{i} = \frac{\psi_{i_{1}} + \psi_{i_{2}} + \psi_{i_{3}}}{3}$$
(4.31)

Noise in the Scaled Surface Normal Element

Given ten images $s^{t}(x,y)$ where t=1-10 of a scaled surface normal element estimate which correspond to a given illumination configuration, a mean temporal noise value is estimated as follows:

$$\psi_{s} = \frac{1}{m} \sum_{x=1}^{\sqrt{m}} \sum_{y=1}^{\sqrt{m}} \left(\frac{1}{9} \sum_{t=1}^{10} (s^{t} - \overline{s}(x, y))^{2} \right)$$
(4.32)

where *m* is the number of pixels in a square image. Since there are three scaled surface normal elements, three temporal noise values are generated $\psi_{s_x}, \psi_{s_y}, \psi_{s_z}$.

Figure of Merit

Empirical estimates of the figures of merit for rough and smooth surfaces M_{rough} and M_{smooth} are calculated with Equations 4.24 and 4.28 using the temporal noise estimates $\overline{\psi}_i, \psi_{s_x}, \psi_{s_y}, \psi_{s_z}$.

4.5 Investigation into Optimal Performance

The figures of merit enabled us to undertake a thorough investigation into the performance of the three-image photometric stereo technique with a view to determining the optimal illumination configuration. We estimated theoretical and empirical values for M_{rough} and M_{smooth} for a series of illumination configurations. In one type of experiment we varied the illumination tilt angle corresponding to the third input image. In a second type of experiment we varied the illumination slant angle common to all three images. The results are presented graphically in Section 4.5.1 and provide a comparison between the theoretical and empirical approaches. In Section 4.5.2 we consider the minimisation of the theoretical expressions.

4.5.1 Graphical Representation

Here we consider three input images which have corresponding illumination tilt angles of τ_1 , τ_2 , τ_3 and a common slant angle σ .

Tilt Angle τ₃ Variation

In these experiments the tilt angles τ_1 and τ_2 and the common slant angle σ were held constant. Their values were chosen to correspond to illumination configurations which are typically employed in photometric stereo. In one set of experiments we used $\tau_1 = 0^\circ$, $\tau_2 = 90^\circ$, $\sigma = 45^\circ$ and in another we used $\tau_1 = 0^\circ$, $\tau_2 = 120^\circ$, $\sigma = 45^\circ$. Both sets of experiments involved altering the illumination configuration by varying the third tilt angle τ_3 (see Figure 4.4). With regard to the images of the real texture, its value was increased by 10° increments over a complete tilt angle rotation ($0^\circ < \tau_3 < 360^\circ$, $\Delta \tau_3 =$ 10°) since images corresponding to these illumination directions are contained in the database. With regard to the theoretical approach we used increments of $1^\circ (\Delta \tau_3 = 1^\circ)$.



Figure 4.4 Example of range of illumination configurations for two tilt angle experiments (Plan view). Increments are $\Delta \tau_3 = 1^{\circ}$ (theoretical), $\Delta \tau_3 = 10^{\circ}$ (empirical).

Figure of merit values were estimated for each configuration. Typical plots are given in Figure 4.5 and Figure 4.6. Plots for M_{smooth} have a similar profile to those for M_{rough} and are presented in Appendix D.



Figure 4.5 Figure of merit M_{rough} versus third tilt angle τ_3 with $\tau_1=0^\circ$, $\tau_2=120^\circ$.



Figure 4.6 Figure of merit M_{rough} versus third tilt angle τ_3 with $\tau_1=0^\circ$, $\tau_2=90^\circ$.

A noticeable feature common to both graphs is that the noise ratio goes off the scale as the third tilt angle coincides with values corresponding to the first and second angles. This is the co-planar situation when the inverse of the illumination matrix does not exist. In this instance it is not possible to solve the system of equations for the unknowns. It is also apparent that increases in the value of the figure of merit become more significant as this situation is approached.

The most interesting feature common to both graphs is that there exists a third tilt angle which corresponds to a minimum. This is approximately 240° when the first and

second tilt angles are set to 0° and 120° respectively as highlighted on the plot. However, if these angles are changed to 0° and 90° the optimal third tilt angle is not 180° but around 225°. This means that McGunnigle's photometric scheme [McGunnigle1998] is sub-optimal but not significantly so and has the advantage of being straightforward to solve.

It is noted that the graphs presented in this section correspond to rough surface textures. Those for smooth surfaces exhibit equivalent behaviour and are given in Appendix D.

Slant Angle σ Variation

In these experiments the three tilt angles τ_1 , τ_2 , τ_3 were held constant. In the set of experiments presented here we used the optimal values $\tau_1 = 0^\circ$, $\tau_2 = 120^\circ$, $\tau_3 = 240^\circ$ determined for a common illumination slant angle configuration. The experiments involved altering the illumination configuration by varying the value of the common slant angle σ . With regard to the images of the real texture, its value was increased in increments of 5° for a range of slant angles ($20^\circ \le \sigma \le 70^\circ$ with $\Delta \sigma = 5^\circ$). With regard to the theoretical approach we used increments of 1° ($\Delta \sigma = 1^\circ$). Figure of merit values were estimated for each configuration.



Figure 4.7 Illustration of range of illumination slant angles in experiments. Increments are $\Delta \sigma = 1^{\circ}$ (*theoretical*), $\Delta \sigma = 5^{\circ}$ (*empirical*).

In this case it is actually the difference in behaviour between the two kinds of surfaces which is interesting. Plots of the figures of merit for rough and smooth surfaces are presented in Figure 4.8 and Figure 4.9. Figure 4.8 demonstrates that with regard to minimising our figure of merit for a texture of rough surface, a slant angle of about 55° is optimal. However, different behaviour is observed for a smooth surface (Figure 4.9). The minimum no longer corresponds to 55° but has increased beyond the range of the

graph. Extrapolation appears to suggest that in this case a slant angle of 90° is optimal. This observation will later be confirmed by minimisation (see Section 4.5.2).



Figure 4.8 Total noise ratio M_{rough} versus slant angle σ with $\tau_1=0^\circ$, $\tau_2=120^\circ$, $\tau_3=240^\circ$.



Figure 4.9 Total noise ratio M_{smooth} versus slant angle σ with $\tau_1=0^\circ$, $\tau_2=120^\circ$, $\tau_3=240^\circ$.

The results of another set of slant angle experiments in which we used illumination tilt angle values of $\tau_1 = 0^\circ$, $\tau_2 = 90^\circ$, $\tau_3 = 180^\circ$ are presented in Appendix D. Similar profiles are apparent with regard to each figure of merit.

4.5.2 Minimisation

The plot profiles given in the previous section each indicate a minimum noise ratio with regard to both tilt and slant angle. An optimisation procedure was used to precisely determine the corresponding parameter values. The figure of merit formulas were minimised by application of the Nelder-Mead algorithm [Press1988, Chap.10]. The algorithm uses a geometrical figure termed a Simplex. For an N-dimension minimisation problem the Simplex is constructed from N+1 points with interconnecting lines. The function is calculated at each of the points and depending on the resulting values, the Simplex is reflected, expanded or contracted with a view to progressing towards the minimum in an iterative manner. This approach facilitated an investigation into the existence of a global minimum.

Three Parameter Minimisation

With a common fixed slant angle the minimum value of the figure of merit was not found to correspond to unique values for the three tilt angles (See Table 4.1). Their values were found to depend on the initial conditions specified. However, it is apparent that the minimum does correspond to a unique *difference* in tilt angle of 120°. This is true for both rough and smooth surfaces. This agrees with the observation from Figure 4.5.

Mrough	$(au_{I,}$	$ au_{2,}$	$ au_3)$
3.1	128.7°	248.7°	8.7°
3.1	0.0°	120.0°	240.0°
3.1	25.3°	145.3°	265.3°

Table 4.1Examples of three parameter minimisation results.

Since our equipment for capturing the images of real textures employs a single moveable light source (see Chapter 3), it is convenient to collect data with a common slant angle. This implies that this result is potentially useful on a practical basis.

Four Parameter Minimisation

(a)

(b)

In this case the slant angle is common for the three illumination vectors as before but its value is not fixed. The minimisation procedure therefore yields the value of four parameters, the three tilt angles and that for the slant (see Table 4.2). Again it is clear that the minimum corresponds to a unique difference in tilt angle values of 120°. This is true for both rough and smooth surfaces. With regard to the slant angle, a unique value of approximately 54.7° is apparent for rough surfaces (see Table 4.2a). This value increases to 90° for smooth surfaces (see Table 4.2b) although this result is not of practical value since in reality light from a source in this position would not impinge on the surface.

<i>M</i> _{rough}	$(au_{l,i})$	$ au_{2,}$	$ au_{3,}$	σ)
3.0	0.0°	120.0°	240.0°	54.7°
3.0	176.4	56.4°	296.4°	54.7°
3.0	324.4°	84.4°	204.4°	54.7°
M_{smooth}	$(\tau_{I_{\perp}})$	$ au_{2}$	$ au_{3}$	<i>σ</i>)
1.6	0.0°	120.0°	240.0°	90.0°
1.6	272.1	32.1°	152.1°	90.0°
1.6	63.8°	303.8°	183.8°	90.0°

Table 4.2Examples of four parameter minimisation results for (a) a rough surface,(b) a smooth surface.

The marked difference in the optimal slant angle for rough and smooth surfaces cannot be attributed to shadowing since our approach does not taken it into account. The results can be explained by considering the difference between the two figures of merit i.e. the noise ratio for the z-component of the scaled surface normal. In this simplified case of common slant angle, the denominator of the noise ratio of the z-component contains $\cos \sigma$ whilst that for the x and y components is $\sin \sigma$ (see Appendix B). Since the z-component is omitted for a smooth surface, a maximum slant results in a minimum M_{smooth} .

We note that the angles determined for the rough surface in Table 4.2a mean that the corresponding optimal illumination vectors are *orthogonal*.

Six Parameter Minimisation

Finally the minimisation was performed in a completely unconstrained manner such that the slant angles were no longer required to be common to the three illumination directions. The examples given in Table 4.3 demonstrate that the conditions corresponding to the minimum noise ratio are not in fact unique but depend on the initial conditions specified in every optimisation. However, in each case the three resulting illumination vectors are *orthogonal*.

M _{rough}	$(\tau_{1,}$	$\tau_{2,}$	$\tau_{3,}$	$\sigma_{1,}$	σ _{2,}	σ3)
3.0	0.0°	112.4°	239.4°	55.3°	56.1°	48.2°
	13.9°	135.0°	255.6°	55.3°	53.3°	55.6°
	70.5°	189.4°	315.3°	59.4°	50.8°	54.3°

Table 4.3Examples of six parameter minimisation results for a rough surface.

To confirm this result the figure of merit M_{rough} was generated over an extensive range of illumination configurations. We produced a 3D scatter plot by plotting M_{rough} against the dot products between the illumination vectors (see Figure 4.10). The minimum of the data cloud corresponds to a dot product of zero between the first and second illumination vectors and a dot product of zero between the second and third illumination vectors. This helps to illustrate the fact that the minimum value for the figure of merit and hence optimal performance corresponds to an orthogonal illumination configuration.



Figure 4.10 M_{rough} versus $\mathbf{b_1}$. $\mathbf{b_2}$ and $\mathbf{b_2}$. $\mathbf{b_3}$.

4.5.3 Summary of Findings

With regard to three-image Lambertian photometric stereo we found that the optimal illumination configuration cannot be specified in terms of a unique set of values for the tilt and slant angles defining illumination direction. Instead we determined that the optimal operating conditions correspond to an orthogonal configuration i.e. when the three illumination vectors are at an angle of 90° to each other (Figure 4.11). This is an interesting result which is somewhat intuitive in hindsight. An orthogonal configuration is not just plausible for isotropic textures. In the case of a unidirectional texture the configuration could be oriented such that one illumination vector points along the grain and has a large slant angle whilst the other two vectors point across the grain with smaller slant angles.



Figure 4.11 Representation of orthogonal vectors

The use of an orthogonal configuration may not be practicable unless the illumination slant angles are constrained to take a common value. In this case the use of a 120° difference in tilt angle is to be recommended. This was found to be applicable to both rough and smooth surfaces. When shadowing is not an issue and the surface is rough in character, a slant angle of around 55° can be used to attain optimal operating conditions. If shadowing is present then this value should be reduced. If on the other hand the surface can be considered to be smooth in nature and not susceptible to shadows then this value can be increased. In this case the term 'grazing' can be used to describe the resulting illumination conditions. It is noted that such textures of this nature will probably not be of interest in visualisation applications since their appearance will not change significantly with different illuminant position although this is dependent on the scale with which the surface is viewed.

4.6 Practical Assessment

The recommendations for optimal placement of the lights presented in the previous section are effectively theoretical. Although an equivalent empirical procedure was carried out to verify the results, its use of images of a real texture with close to ideal reflectance properties meant that the validity of the proposed illumination conditions was uncertain for the non-ideal case. A quantitative assessment of a practical nature was therefore undertaken to investigate this issue. This involved the use of the thirty-one real textures described in Chapter 3 many of which are prone to non-ideal reflectance characteristics such as shadowing and specular highlights. Images of the textures corresponding to a variety of illumination configurations were processed using the three-image photometric stereo algorithm. The Lambertian model (Equation 2.2) was then used to relight the generated surface gradient and albedo images. The relit images produced were then compared to the originals. The comparison was not just visual; the difference between them was quantitatively measured in terms of the texture signal to relight error ratio *TSER* (see Equation 3.10).

Tilt Angle Spacing

Three different tilt angle spacings were utilised in the experiments. These correspond to the theoretical optimal ($\Delta \tau$ =120°), McGunnigle's simplified photometric scheme ($\Delta \tau$ =90°) and finally an asymmetric arrangement ($\Delta \tau$ =50°). A constant slant angle of 45° was used in every case. Following each application of the algorithm relit images were generated over a complete revolution in terms of tilt angle ($\Delta \tau$ =10°) for the two other slant angle values present in the image database i.e. σ =30° and 60°. This avoided the case when the illumination conditions of a relit image correspond to that of an input image when the relight error would tend towards zero. For each texture the signal to relight error ratio *SER* was estimated for each of the 72 illumination directions using Equation 3.9. A value for the texture signal to relight error ratio *TSER* which we employ as a measure of relighting accuracy was then obtained from Equation 3.10. This was calculated for each of the three tilt angle spacings with every texture. The results are presented in the bar chart in Figure 4.12.



Figure 4.12 Practical evaluation of the proposed optimal illumination conditions.

Whilst it is evident that some of the textures are more suited to use with Lambertian photometric stereo than others, it is apparent that the proposed optimal illumination arrangement has outperformed the other arrangements in all but one case. The exception (texture ac) relates to a specular texture with less than ideal reflection. The tilt angle spacing of 90° performs well but the resulting *TSER* values are relatively low in general compared to the optimal case. The asymmetric case performs poorly in comparison. This is also apparent from the samples of relit images given in Figure 4.13 on comparison with the original image. We highlight this in Figure 4.14 by presenting a plot of the intensities taken from a single row of the difference images. For clarity, we use a moving average to identify the underlying trend. Again it is clear that the optimal configuration provides a consistently better approximation to the original intensity image.

Overall, we have verified the theoretical results presented earlier in the chapter. We conclude that the optimal spacing of the illumination vectors is 120° with regard to tilt angle when they are constrained to be of common slant angle and that there is a benefit in terms of accuracy when this configuration is utilised on a practical basis.

(a) Original image of texture with illumination conditions $\tau = 270^{\circ}$ and $\sigma = 45^{\circ}$.



(b) Optimal conditions -120° spacing with regard to tilt for constant slant.





(c) McGunnigle's scheme -90° spacing with regard to tilt for constant slant.





(d) Asymmetric arrangement - 50° spacing with regard to tilt for constant slant.





Figure 4.13 Comparison of original image with relights (texture n).

(a) Original image of texture with illumination conditions $\tau = 270^{\circ}$ and $\sigma = 45^{\circ}$.



(b) Optimal conditions -120° spacing with regard to tilt for constant slant.





(c) McGunnigle's scheme -90° spacing with regard to tilt for constant slant.





(d) Asymmetric arrangement - 50° spacing with regard to tilt for constant slant.





Figure 4.13 (Cont'd) Comparison of original image with relights (texture q).



Figure 4.14 Moving average (10-pixel) trend for a single row of the difference images given in Figure 4.13.

Slant Angle Selection

Experiments concerning the slant angle were also undertaken. In this case illumination configurations using the optimal tilt angle spacing of 120° were used with a range of slant angles common to each illumination vector $(30^{\circ} \le \sigma \le 60^{\circ} \text{ with } \Delta \sigma = 15^{\circ})$. Following each application of the algorithm relit images were generated over a complete revolution in terms of tilt angle ($\Delta \tau = 10^{\circ}$) for the two other slant angle values. For example, if the input images illumination slant angle is 30° then seventy-two relit images would be generated with slant angles of 45° and 60° . Whilst this approach is not ideal it provides some insight into the issue of slant angle selection which is potentially dependent on the surface type according to the theory and is furthermore constrained by the need to minimise the presence of shadows. The results are presented in Figure 4.15.



Figure 4.15 Practical evaluation of effect of slant angle on accuracy based on optimal illumination configuration with regard to tilt angle ($\Delta \tau = 120^{\circ}$).

Assuming that the comparison is valid it is evident that in general it is preferable to use a slant angle of 45° because this intermediate value corresponds to the best performance of photometric stereo for the majority of real textures. This result merely helps to confirm that the effect of shadowing is important. It is not really possible to distinguish between our proposed optimal value of around 55° for rough surface texture and Woodham's recommendation of maximising slant angle because of the impact of shadowing for larger slant angles. Shadowing is not taken into consideration in either of the theoretical evaluations. However, based on evidence from these practical results for thirty-one textures it is prudent to avoid the use of extreme slant values. Overall a value of 45° appears to be more appropriate for textures of a similar nature to that used in the investigation.

4.7 Summary & Discussion

The results of an investigation into the optimal placement of the illumination vectors in the three-image Lambertian photometric stereo technique were reported in this chapter. The work is based on an overall figure of merit which was derived from noise variance expressions. This metric was developed by considering the image-based rendering of a Lambertian surface. The resulting equation is in terms of the illumination tilt and slant angles defining the orientation of the three illumination vectors and it was employed to gauge the effect of various illumination configurations. The optimal configuration was determined by its minimisation. The theoretical results were confirmed through the use of an equivalent empirical method.

The optimal illumination configuration was determined to correspond to an orthogonal arrangement of the three illumination vectors. With regard to a more practical implementation of the technique when the slant angle is common for the three vectors, we found the optimal difference between successive tilt angles to be 120°. The advantage of using this optimal lighting arrangement was verified by assessing the relighting error of thirty-one real textures. Relative to the other illumination arrangements tested a benefit in terms of error reduction is evident in almost every case.

Using an illumination configuration of $\Delta \tau = 120^{\circ}$ and common slant angle, the derived figure of merit was used to show that optimal performance corresponds to a maximum slant angle of 90° when the texture is smooth. This is the ideal case. In reality this would not be possible due to shadowing considerations. This finding is equivalent to Woodham's original observation that accuracy in photometric stereo is improved if a large slant angle is used [Woodham1980]. He based his argument on the fact that increasing the slant angle increases the density of the corresponding reflectance map. This is desirable because a large change in intensity will result from a small change in the surface gradients *p* and *q*. In other words Woodham recommends that the sensitivity of the intensity with respect to the surface gradients is maximised. Our approach to this problem is similar but 'inverted' since we tackle the issue by minimising the sensitivity of the scaled surface normal with respect to the intensity.

For a rough surface illuminated under this configuration i.e. $\Delta \tau = 120^{\circ}$ and common slant angle we found the optimal slant angle to be 55°. This agrees with the more general case when the illumination vectors are orthogonal. Although Woodham advocates the use of a large slant angle he does qualify this statement by indicating that a compromise exists since shadowing should be minimised [Woodham1980]. This

would certainly be an issue for a rough surface texture. However, Woodham does not actually differentiate between types of surfaces and we acknowledge the divergence in results at this point. We attribute this to the fact that the scaled surface normal sensitivity was minimised in our approach rather than the surface gradients.

Whilst the optimal illumination conditions proposed in this chapter are potentially useful on a practical basis, it is noted that they have been derived for a texture exhibiting ideal diffuse reflection. The fact that neither shadows, specularities nor interreflections have been considered in the development of the theory means that the application of such guidelines should really be restricted. However, the assessment of the three-image photometric stereo technique with real textures demonstrated that even in the presence of shadows and specularities, using the 120° tilt angle spacing with constant slant angle was in fact relatively beneficial. In the case of specular reflectance this may well be because the specular peak is narrow and therefore not frequently observed under the three light positions. For textures which exhibit far from ideal reflectance the recommended illumination arrangement can simply be used as a first guess of optimal illumination conditions.

4.8 Conclusions

Overall, we conclude that a difference between successive tilt angles of 120° is to be recommended when the illumination configuration is constrained to have a common slant angle. Based on the theoretical and empirical evidence presented we recommend the use of a maximum slant angle of 55° for rough surface textures. These observations have potential implications for three-image photometric stereo algorithms which select the three 'best' pixel values from multiple images [Rushmeier1997, Coleman1982, Petrou2001]. Once the shadowed and specular intensities have been discarded, it is very possible that the illumination conditions corresponding to the remaining intensities may be less than optimal.

Chapter 5

Uncalibrated Approach to Photometric Stereo

5.1 Introduction

The ability to implement an *uncalibrated* technique would be of great value in obtaining surface representations of textures. Lessening the burden of input information, which implicitly involves calibration and measurement, would make the procedure much more amenable to practical implementation. Hence it is a goal worth pursuing provided that accuracy is not compromised. In this chapter we consider ways of achieving this objective. We are specifically concerned with the case when image intensity is the only input unlike calibrated photometric stereo where illumination direction information pertaining to each image is also available. Having reviewed the available techniques and discussed the issues involved with each, the ultimate goal of this chapter is then to identify those methods which are most suited to our specific application of determining surface representations of textured surfaces.

This chapter is organised as follows :

The primary step in the uncalibrated photometric stereo method is introduced in Section 5.2. The main issue which concerns this technique is then discussed and the various means of tackling it proposed in the literature are described in detail. In Section 5.3 a summary of the various techniques is given and we conclude by proposing the most suitable for use in our specific case.

5.2 Uncalibrated Photometric Stereo

Hayakawa [Hayakawa1994] proposed a photometric stereo method for estimating surface representations of objects without prior knowledge of the illumination. He assumes Lambertian behaviour such that the image intensity data matrix I can be written as the product of two matrices, S & L. These represent the true scaled surface normals and scaled illumination vectors respectively, neither of which are known in this case and are to be determined. This was originally introduced in Chapter 2 but we re-state the equations here for clarity.

$$\mathbf{I} = \mathbf{SL} \tag{2.35}$$

It is possible to factorise the input intensity matrix into a *pseudo* surface matrix and a *pseudo* illumination matrix, $\hat{S} \& \hat{L}$. These represent a possible solution but it is not unique and an ambiguity exists since the following holds.

$$\mathbf{I} = \hat{\mathbf{S}}\hat{\mathbf{L}} = \hat{\mathbf{S}}\mathbf{A}^{\mathrm{T}}\mathbf{A}^{-\mathrm{T}}\hat{\mathbf{L}}$$
(2.39)

This is a specific example of the generic bilinear calibration-estimation problem described by Koenderink [Koenderink1997]. Even though the initial decomposition to obtain the first estimates is straightforward in terms of mathematics, the determination of the ambiguity matrix is more challenging.

5.2.1 Ambiguity Reduction/Resolution Survey

The ambiguity **A** is an arbitrary 3×3 matrix. It describes a general linear transformation and is written in group notation as GL(3). It is noted that its determinant cannot be zero because it must be invertible.

$$\mathbf{A} = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix} \qquad \text{where } a_{ij} \in \mathbb{R}$$

$$(5.1)$$

In effect nine degrees of freedom must be resolved in order to derive a unique solution for the surface representation and illumination estimates.

$$\mathbf{S} = \mathbf{\hat{S}}\mathbf{A}^{\mathrm{T}} \qquad \mathbf{L} = \mathbf{A}^{-\mathrm{T}}\mathbf{\hat{L}} \qquad (5.2, 5.3)$$

The ambiguity can only be determined through knowledge of surface and illumination characteristics. The approaches detailed in the literature mostly centre on identifying ways of satisfying this requirement without involving calibration. The alternative sources of information and assumptions proposed will be considered in the following subsections and their suitability for our texture-specific application will be assessed. Those methods which utilise extra equipment such as calibration objects to provide knowledge of the object class [Yuille1999] or polarisers [Drbohlav2001] will not be considered.

We emphasise here that Sections 5.2.1.1-5.2.1.3 are essentially a more detailed literature survey of uncalibrated photometric stereo. In this case the specific mathematics involved in the development of each algorithm is presented. The aim is to provide sufficient detail for readers who intend to conduct research in the area. Other readers may wish to proceed directly to the summary in Section 5.3.

5.2.1.1 Vector Magnitude Constraint

Hayakawa detailed two constraints which facilitate the determination of the ambiguity [Hayakawa1994]. He advocates finding either six pixels in which the albedo is known or constant or alternatively, six frames in which the illuminant intensity is known or constant. The treatment for both constraints is equivalent since it basically involves analysing six scaled vectors from the pseudo estimates which should be of the same length. The constant albedo constraint will be considered in the following.

In this case we focus on the scaled surface normal vectors. According to Equation 5.2, a true estimate of a surface normal can be written as the product of its corresponding pseudo estimate and the ambiguity matrix.

$$\mathbf{A}\begin{bmatrix}\hat{s}_{x}\\\hat{s}_{y}\\\hat{s}_{z}\end{bmatrix} = \begin{bmatrix}s_{x}\\s_{y}\\s_{z}\end{bmatrix}$$
(5.4)

It is noted that the magnitude of the scaled surface normal is the albedo.

$$|\mathbf{A}\hat{\mathbf{s}}| = |\mathbf{s}| = \rho \tag{5.5}$$

The dot product of the vector with itself therefore results in the following:

$$\hat{\mathbf{s}}^{\mathrm{T}}\mathbf{A}^{\mathrm{T}}\mathbf{A}\ \hat{\mathbf{s}} = \rho^2 \tag{5.6}$$

For convenience we define a 3×3 matrix **C** as follows:

$$\mathbf{C} = \mathbf{A}^{\mathrm{T}} \mathbf{A} \tag{5.7}$$

C is a symmetric matrix as can be seen if written explicitly in terms of the elements of A.

$$\mathbf{C} = \begin{bmatrix} a_{11}^2 + a_{21}^2 + a_{31}^2 & a_{11}a_{12} + a_{21}a_{22} + a_{31}a_{32} & a_{11}a_{13} + a_{21}a_{23} + a_{31}a_{33} \\ a_{11}a_{12} + a_{21}a_{22} + a_{31}a_{32} & a_{12}^2 + a_{22}^2 + a_{32}^2 & a_{12}a_{13} + a_{22}a_{23} + a_{32}a_{33} \\ a_{11}a_{13} + a_{21}a_{23} + a_{31}a_{33} & a_{12}a_{13} + a_{22}a_{23} + a_{32}a_{33} & a_{13}^2 + a_{23}^2 + a_{33}^2 \end{bmatrix}$$

$$= \begin{bmatrix} c_1 & c_2 & c_3 \\ c_2 & c_4 & c_5 \\ c_3 & c_5 & c_6 \end{bmatrix}$$
(5.8)

Inserting this symmetric matrix into Equation 5.6 and expanding, the following is obtained. $2^{T} G^{2} = 2^{2}$

Since C is symmetric it contains six independent parameters. If six pseudo scaled surface vectors with constant or known albedo can be identified then the following system of linear equations (based on Equation 5.9) can be solved by inverting the matrix of knowns i.e. pseudo vectors to determine C.

$$\begin{bmatrix} \hat{s}_{1,x}^{2} & 2\hat{s}_{1,x}\hat{s}_{1,y} & 2\hat{s}_{1,x}\hat{s}_{1,z} & \hat{s}_{1,y}^{2} & 2\hat{s}_{1,y}\hat{s}_{1,z} & \hat{s}_{1,z}^{2} \\ \hat{s}_{2,x}^{2} & 2\hat{s}_{2,x}\hat{s}_{2,y} & 2\hat{s}_{2,x}\hat{s}_{2,z} & \hat{s}_{2,y}^{2} & 2\hat{s}_{2,y}\hat{s}_{2,z} & \hat{s}_{2,z}^{2} \\ \hat{s}_{3,x}^{2} & 2\hat{s}_{3,x}\hat{s}_{3,y} & 2\hat{s}_{3,x}\hat{s}_{3,z} & \hat{s}_{3,y}^{2} & 2\hat{s}_{3,y}\hat{s}_{3,z} & \hat{s}_{3,z}^{2} \\ \hat{s}_{4,x}^{2} & 2\hat{s}_{4,x}\hat{s}_{4,y} & 2\hat{s}_{4,x}\hat{s}_{4,z} & \hat{s}_{4,y}^{2} & 2\hat{s}_{4,y}\hat{s}_{4,z} & \hat{s}_{4,z}^{2} \\ \hat{s}_{5,x}^{2} & 2\hat{s}_{5,x}\hat{s}_{5,y} & 2\hat{s}_{5,x}\hat{s}_{5,z} & \hat{s}_{5,y}^{2} & 2\hat{s}_{5,y}\hat{s}_{5,z} & \hat{s}_{5,z}^{2} \\ \hat{s}_{6,x}^{2} & 2\hat{s}_{6,x}\hat{s}_{6,y} & 2\hat{s}_{6,x}\hat{s}_{6,z} & \hat{s}_{6,y}^{2} & 2\hat{s}_{6,y}\hat{s}_{6,z} & \hat{s}_{6,z}^{2} \\ \end{bmatrix} \begin{bmatrix} \rho^{2} \\ \rho^{2} \\ \rho^{2} \\ \rho^{2} \\ \rho^{2} \\ \rho^{2} \\ \rho^{2} \end{bmatrix}$$
(5.10)

In the case when the albedo value is constant but unknown, it is assigned the value of unity. The implication is that a constant scaling factor is introduced into the solution.

$$\mathbf{C} = k_1 \mathbf{C}_{true} \quad \text{where} \quad k_1 = \frac{1}{\rho^2} \tag{5.11}$$

The resulting symmetric matrix C can then be factorised to obtain an estimate of the ambiguity.

Residual Ambiguity

Whilst C is determined uniquely this is not the case for the resulting ambiguity A which has been obtained by factorisation and a residual ambiguity R exists. The form of the residual ambiguity is considered in the following.

$$\mathbf{C} = \mathbf{A}^{\mathrm{T}}\mathbf{A} = (\mathbf{R}\mathbf{A})^{\mathrm{T}}\mathbf{R}\mathbf{A}$$
$$= \mathbf{A}^{\mathrm{T}}\mathbf{R}^{\mathrm{T}}\mathbf{R}\mathbf{A}$$
(5.12)

For this to hold, the product of the residual ambiguity and its transpose must equate to the identity matrix. This implies that the transpose is equal to the inverse. The residual ambiguity therefore corresponds to an orthogonal transformation. In group notation this is written as O(3) and corresponds to three degrees of freedom.

$$\mathbf{R}^{\mathrm{T}}\mathbf{R} = \mathbf{I}_{d}$$

$$\downarrow \qquad (5.13)$$

$$\mathbf{R}^{\mathrm{T}} = \mathbf{R}^{-1}$$

Overall this constant albedo constraint reduces the ambiguity from a general linear transformation to an orthogonal one. The remaining ambiguity is therefore equivalent to a rotation and corresponds to 3 degrees of freedom. If the constant albedo is assumed to be unity then there is also an extra scaling factor to consider with regard to fully resolving the ambiguity.

5.2.1.2 Integrability Constraint

An integrability constraint is a further option with regard to ambiguity reduction if the surface in question can be considered to be smooth and continuous. This theory was originally developed by Belhumeur et al [Belhumeur1999] although the following is based on more detailed explanations given in [Drbohlav2003] and [Georghiades2003a].

In this case the second partial derivatives of the surface height are equated. This can be written in terms of surface gradients as follows:

$$\frac{\partial p}{\partial y} = \frac{\partial q}{\partial x} \tag{5.14}$$

Given a scaled surface normal $\mathbf{s} = [s_x \ s_y \ s_z]^T$, Equation 5.14 can be re-stated and then expanded as follows.

$$\frac{\partial}{\partial y} \left(\frac{s_x}{s_z} \right) = \frac{\partial}{\partial x} \left(\frac{s_y}{s_z} \right) \xrightarrow{\frac{\partial}{\partial x} \left(\frac{u}{v} \right) = \frac{v \frac{\partial x}{\partial x} - u \frac{\partial x}{\partial x}}{v^2}} s_z \frac{\partial s_y}{\partial x} - s_y \frac{\partial s_z}{\partial x} = s_z \frac{\partial s_x}{\partial y} - s_x \frac{\partial s_z}{\partial y}$$
(5.15)

If the relationship between a true scaled surface normal and its pseudo counterpart given by Equation 5.4 is again considered, Equation 5.15 can be developed in terms of the pseudo normals and the ambiguity matrix. Individual rows of the ambiguity matrix are defined as $\mathbf{a}_i = \begin{bmatrix} a_{i1} & a_{i2} & a_{i3} \end{bmatrix}$ where i=1-3 to facilitate this. The scaled surface normal is now written as $\mathbf{s} = \begin{bmatrix} \mathbf{a}_{1} \hat{\mathbf{s}} & \mathbf{a}_{2} \hat{\mathbf{s}} & \mathbf{a}_{3} \hat{\mathbf{s}} \end{bmatrix}^{\mathrm{T}}$. Inserting its elements into Equation 5.15 the following is obtained.

$$\mathbf{a}_{3}\hat{\mathbf{s}}\frac{\partial(\mathbf{a}_{2}\hat{\mathbf{s}})}{\partial x} - \mathbf{a}_{2}\hat{\mathbf{s}}\frac{\partial(\mathbf{a}_{3}\hat{\mathbf{s}})}{\partial x} = \mathbf{a}_{3}\hat{\mathbf{s}}\frac{\partial(\mathbf{a}_{1}\hat{\mathbf{s}})}{\partial y} - \mathbf{a}_{1}\hat{\mathbf{s}}\frac{\partial(\mathbf{a}_{3}\hat{\mathbf{s}})}{\partial x}$$
(5.16)

This is re-arranged to provide a single constraint for each surface facet.

$$\hat{\mathbf{s}}^{\mathrm{T}}\mathbf{H}_{1}\frac{\partial\hat{\mathbf{s}}}{\partial x} - \hat{\mathbf{s}}^{\mathrm{T}}\mathbf{H}_{2}\frac{\partial\hat{\mathbf{s}}}{\partial y} = 0$$
(5.17)

where
$$\mathbf{H}_1 = \mathbf{a}_3^T \mathbf{a}_2 - \mathbf{a}_2^T \mathbf{a}_3$$
 & $\mathbf{H}_2 = \mathbf{a}_3^T \mathbf{a}_1 - \mathbf{a}_1^T \mathbf{a}_3$

Equation 5.17 is linear in the elements of the matrices $\mathbf{H}_1 \& \mathbf{H}_2$ which are skew symmetric and contain six independent parameters between them.

$$\mathbf{H}_{1} = \begin{bmatrix} 0 & -h_{1} & h_{3} \\ h_{1} & 0 & -h_{2} \\ -h_{3} & h_{2} & 0 \end{bmatrix} \qquad \mathbf{H}_{2} = \begin{bmatrix} 0 & -h_{4} & h_{6} \\ h_{4} & 0 & -h_{5} \\ -h_{6} & h_{5} & 0 \end{bmatrix}$$
(5.18, 5.19)

where

$$\mathbf{h} = \begin{bmatrix} h_1 \\ h_2 \\ h_3 \\ h_4 \\ h_5 \\ h_6 \end{bmatrix} = \begin{bmatrix} a_{32}a_{21} - a_{31}a_{22} \\ a_{33}a_{22} - a_{32}a_{23} \\ a_{31}a_{23} - a_{33}a_{21} \\ a_{32}a_{11} - a_{31}a_{12} \\ a_{33}a_{12} - a_{32}a_{13} \\ a_{31}a_{13} - a_{33}a_{11} \end{bmatrix}$$
(5.20)

With regard to solving for the elements of \mathbf{h} , it is useful to define the following vector for the scaled surface normal \mathbf{s} of each facet $i_{.}$

$$\mathbf{g}_{i} = \begin{bmatrix} \hat{s}_{i,y} \frac{\partial \hat{s}_{i,x}}{\partial x} - \hat{s}_{x} \frac{\partial \hat{s}_{i,y}}{\partial x} \\ \hat{s}_{i,z} \frac{\partial \hat{s}_{i,y}}{\partial x} - \hat{s}_{i,y} \frac{\partial \hat{s}_{i,z}}{\partial x} \\ \hat{s}_{i,x} \frac{\partial \hat{s}_{i,z}}{\partial x} - \hat{s}_{i,z} \frac{\partial \hat{s}_{i,x}}{\partial x} \\ \hat{s}_{i,x} \frac{\partial \hat{s}_{i,z}}{\partial y} - \hat{s}_{i,y} \frac{\partial \hat{s}_{i,x}}{\partial y} \\ \hat{s}_{i,y} \frac{\partial \hat{s}_{i,z}}{\partial y} - \hat{s}_{i,z} \frac{\partial \hat{s}_{i,y}}{\partial y} \\ \hat{s}_{i,z} \frac{\partial \hat{s}_{i,z}}{\partial y} - \hat{s}_{i,z} \frac{\partial \hat{s}_{i,z}}{\partial y} \end{bmatrix}$$

$$(5.21)$$

Equation 5.17 can hence be re-stated in the following form :

$$\mathbf{g}_i^{\mathrm{T}}\mathbf{h} = 0 \tag{5.22}$$

This provides a total of m constraints since there are a corresponding number of surface facets. These constraints can be written together in matrix form as follows :

$$\mathbf{Gh} = \mathbf{0}$$
where $\mathbf{G} = [\mathbf{g}_1^{\mathsf{T}} \ \mathbf{g}_2^{\mathsf{T}} \dots \mathbf{g}_m^{\mathsf{T}}]^{\mathsf{T}}$
(5.23)

This over-constrained system of linear equations can be solved for \mathbf{h} up to a scale by finding the null space of \mathbf{G} .

The effect of applying the integrability constraint can be gauged by considering the inverse transpose of the original ambiguity matrix [Drbohlav2003]. It is written here in terms of the matrix adjugate divided by the determinant.

$$\mathbf{A}^{-\mathbf{T}} = \frac{1}{\det(\mathbf{A})} \begin{bmatrix} a_{33}a_{22} - a_{32}a_{23} & a_{31}a_{23} - a_{33}a_{21} & a_{32}a_{21} - a_{31}a_{22} \\ a_{32}a_{13} - a_{33}a_{12} & a_{33}a_{11} - a_{31}a_{13} & a_{31}a_{12} - a_{32}a_{11} \\ a_{12}a_{23} - a_{13}a_{22} & a_{13}a_{21} - a_{11}a_{23} & a_{11}a_{22} - a_{12}a_{21} \end{bmatrix}$$
$$= \frac{1}{\det(\mathbf{A})} \begin{bmatrix} h_2 & h_3 & h_1 \\ -h_5 & -h_6 & -h_4 \\ a_{31}^{-T} & a_{32}^{-T} & a_{33}^{-T} \end{bmatrix}$$
(5.24)

Since **h** has been determined, the elements of the top two rows are fixed in value whilst the elements in the third row take arbitrary values. As long as the resulting matrix in invertible then the corresponding matrix \mathbf{A} which ensures that the transformed scaled surface normals are integrable can be determined.

Residual Ambiguity

The ambiguity has not been resolved completely by using the integrability constraint due to the requirement of arbitrary values and it is apparent that a residual ambiguity **R** exists. In terms of the inverse transpose, it transforms $\mathbf{A}^{-\mathsf{T}}$ into $\mathbf{R}^{-\mathsf{T}}\mathbf{A}^{-\mathsf{T}}$. Since the integrability constraint must continue to hold, the form of **R** is limited because the top two rows of $\mathbf{A}^{-\mathsf{T}}$ containing the elements of **h** cannot change. The only applicable transformations are :

$$\mathbf{R}^{-\mathrm{T}} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ k_2 & k_3 & k_4 \end{bmatrix} \longrightarrow \mathbf{R} = \begin{bmatrix} 1 & 0 & -\frac{\kappa_2}{k_4} \\ 0 & 1 & -\frac{k_3}{k_4} \\ 0 & 0 & \frac{1}{k_4} \end{bmatrix} \qquad \begin{array}{c} k_2, k_3 \in \mathbb{R} \\ k_4 \neq 0 \end{array}$$
(5.25)

The k_4 parameter must be non-zero for the matrix to be invertible.

Since the elements of **h** are only determined up to a scale, the final form of the residual ambiguity is a group of generalised bas-relief transformations (GBR) [Belhumeur1999].

5.2.1.3 Consistent Viewpoint Constraint

If surface reflectance is composed of both a diffuse and a specular component then the ambiguity can be reduced according to the consistent viewpoint constraint [Drbohlav2002]. The presence of highlights in images of the illuminated surface is exploited to provide geometrical information about the viewing direction. Although we focus on Lambertian photometric stereo in this thesis, mirror-like specular reflectance is assumed by Drbohlav. This means that specular highlights in the images are sparse and can be treated as outliers to the diffuse model. Lambertian photometric stereo can therefore still be utilised and in this case has the advantage of additional information provided by this constraint.



Figure 5.1 Mirror-like specular reflection from a surface facet such that the highlight is visible.

If a highlight is observed in an image of the illuminated surface then the surface normal of the corresponding facet bisects the illumination and viewing vectors as depicted in Figure 5.1. The viewing vector can hence be written as follows :

$$\mathbf{v} = 2 \left(\frac{\mathbf{l}}{|\mathbf{l}|} \cdot \frac{\mathbf{s}}{|\mathbf{s}|} \right) \frac{\mathbf{s}}{|\mathbf{s}|} - \frac{\mathbf{l}}{|\mathbf{l}|}$$
(5.27)

It is important to note that the view vector will be the same for highlights which correspond to other pairs of surface normals and illumination vectors, termed *specular pairs*.

Written in terms of the ambiguity and the pseudo scaled surface normals and illumination vectors attained from factorisation this becomes :

$$\mathbf{v} = \frac{2[(\mathbf{A}^{-T}\hat{\mathbf{l}}) \cdot (\mathbf{A}\hat{\mathbf{s}})]\mathbf{A}\hat{\mathbf{s}}}{|\mathbf{A}^{-T}\hat{\mathbf{l}}| |\mathbf{A}\hat{\mathbf{s}}|^2} - \frac{\mathbf{A}^{-T}\hat{\mathbf{l}}}{|\mathbf{A}^{-T}\hat{\mathbf{l}}|}$$
(5.28)

Dividing through by $|\mathbf{A}^{-T}\hat{\mathbf{l}}| |\mathbf{A}\hat{\mathbf{s}}|^2 \mathbf{A}^{T}$ and then making substitutions the following is obtained.

$$\alpha(\hat{\mathbf{l}}, \hat{\mathbf{s}})\mathbf{w} = \mathbf{2}(\hat{\mathbf{l}} \cdot \hat{\mathbf{s}})\mathbf{C}\hat{\mathbf{s}} - (\hat{\mathbf{s}} \cdot \mathbf{C}\hat{\mathbf{s}})\hat{\mathbf{l}}$$
(5.29)

where $\alpha(\hat{\mathbf{l}}, \hat{\mathbf{s}}) = |\mathbf{A}^{-T}\hat{\mathbf{l}}| ||\mathbf{A}\hat{\mathbf{s}}|^2 ||\mathbf{A}^{T}|$,

$$\mathbf{w} = \frac{\mathbf{A}^{\mathrm{T}}\mathbf{v}}{|\mathbf{A}^{\mathrm{T}}\mathbf{v}|},$$
$$\mathbf{C} = \mathbf{A}^{\mathrm{T}}\mathbf{A}.$$

Equation 5.29 represents a system of non-linear equations if at least four specular pairs can be identified; they are selected by hand from regions of highlights in the input images which have been segmented by thresholding. The objective is to solve this in order to determine the transformation **A** which maps the pseudo normals and vectors onto those which fulfil the consistent viewpoint constraint. An algorithm based on bundle adjustment is used to achieve this. The resulting pseudo viewing vector **w** is then rotated to align it with the true viewing vector $\mathbf{v} = [0, 0, 1]^{T}$.

Residual Ambiguity

For specular pairs which obey the consistent viewpoint constraint and whose viewing vector is \mathbf{v} , \mathbf{C} can only take the form of a scaled identity matrix if Equation 5.29 is to hold such that the pseudo viewing vector $\mathbf{w} = \mathbf{v}$. This allows the form of the residual ambiguity \mathbf{R} to be gauged.

$$\begin{array}{c} \mathbf{C} = k_9^2 \mathbf{I}_d \\ \mathbf{C} = \mathbf{R}^{\mathrm{T}} \mathbf{R} \end{array} \end{array} \right\} \qquad \qquad \mathbf{R} = k_9 \mathbf{O} \\ \text{where } \mathbf{O} \in O(3) \end{array}$$

$$(5.30)$$

Although \mathbf{R} is a scaled orthogonal transformation it is limited further by the fact that the true viewing direction has been fixed. In order to preserve the consistent viewing direction constraint the only possible invertible transformations are :

$$\mathbf{R} = k_9 \begin{bmatrix} \pm 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \mathbf{R}_z$$
(5.31)

The residual ambiguity therefore entails a rotation about the *z*-axis, a reflection about the *z*-*y* plane which concerns the co-ordinate frame handedness and an unknown scaling factor. None of these transformations affect the direction of the viewing vector. Drbohlav resolves these remaining degrees of freedom by employing the integrability constraint.

Georghiades recently presented a generalisation of the consistent viewpoint constraint by incorporating the Torrance-Sparrow model into uncalibrated photometric stereo [Georghiades2003a/b/c]. This physically-based reflectance model consists of a diffuse component and a surface scatter component. The latter enables a specular lobe in the forward direction to be modelled; this is a characteristic of specular reflection from rough surfaces where the highlight is not simply a mirror reflection of the incoming ray of light. In this case the uncalibrated photometric stereo ambiguity is reduced because the specular lobes in the input intensity images must be aligned with those derived from the T-S model.

Georghiades initially applied the integrability constraint to reduce the ambiguity to a GBR residual ambiguity \mathbf{R} as defined in Section 5.2.1.2. The consistent viewpoint constraint was then formulated by considering the image intensity for a surface facet given by the T-S model simplified by making justifiable assumptions.

$$i_{TS} = \mathbf{R}\hat{\mathbf{s}} \cdot \mathbf{R}^{-T}\hat{\mathbf{l}} + \rho_{s} |\mathbf{R}^{-T}\hat{\mathbf{l}}| \exp(-\upsilon^{2}[\cos^{-1}\left(\frac{\mathbf{R}\hat{\mathbf{s}} \cdot (\frac{\mathbf{R}^{-T}\hat{\mathbf{l}}}{|\mathbf{R}^{-T}\hat{\mathbf{l}}|} + \mathbf{v})}{|\mathbf{R}\hat{\mathbf{s}}||\frac{\mathbf{R}^{-T}\hat{\mathbf{l}}}{|\mathbf{R}^{-T}\hat{\mathbf{l}}|} + \mathbf{v}|}\right)]^{2})(\mathbf{v} \cdot \frac{\mathbf{R}\hat{\mathbf{s}}}{|\mathbf{R}\hat{\mathbf{s}}|})^{-1}$$
(5.32)

where v = surface roughness, $\rho_s =$ specular albedo

The GBR ambiguity is resolved and the parameters of the T-S model are determined by iterative least squares techniques using a cost function based on the difference between predicted and actual intensities.

5.3 Summary & Discussion

In this chapter we introduced uncalibrated photometric stereo and reviewed the various approaches utilised to reduce and resolve the inherent ambiguity. A summary is given in the following table in order to provide an overview.

Author	Constraints	Ambiguity Red	Comments	
		Transform	d.o.f	
		GL(3)	9	
		\downarrow	\downarrow	
Hayakawa	vector magnitude	k O (3)	3	6 albedos/light intensities constant
			(+ 1 scale)	or known.
Belhumeur,	integrability		3	Smooth surfaces.
Kriegman & Yuille		GBR		
Belhumeur,	integrability +	$GBR \to k\mathbf{I}_{d\mathbf{t}}$	- -	Smooth surfaces.
Kriegman & Yuille	vector magnitude		$3 \rightarrow \mathbf{I}$ (+ 1 scale)	intensities constant.
Drbohlav, Šára	consistent viewpoint	$k({}^{\pm}\mathbf{I}_{d\pm})\mathbf{R}_{z}$	2 (+ 1 scale)	Ideal specular reflectance.
Drbohlav, Šára	consistent viewpoint + integrability	$k({}^{\pm}I_{d\pm}) R_{z} \rightarrow k I_{d\pm}$	$2 \rightarrow \frac{1}{(+ 1 \text{ scale})}$	Smooth surfaces. Ideal specular reflectance.
Georghiades	integrability + consistent viewpoint	$GBR \rightarrow k I_{d\pm}$	$3 \rightarrow 1$ (+ 1 scale)	Smooth surfaces. Torrance-Sparrow.

where GL(3) is a general linear transform, O(3) is an orthogonal transform,

GBR is generalised bas-relief transform, R_z is a *z*-axis rotation,

 $I_{d\pm}$ is a convex/concave transform,

 $I_{d\pm}$ is a convex/concave & co-ordinate frame handedness transform.

Table 5.1Summary of literature review of uncalibrated photometric stereo with
emphasis on ambiguity reduction/resolution.

It is apparent that each of the approaches detailed above has utilised at least one of three main types of constraints : integrability, consistent viewpoint & constant vector magnitude. Whilst these are potential candidates for ambiguity resolution, it is evident that their application is very much dependent on the characteristics of the surface in question. We have chosen rough surface textures which largely exhibit diffuse reflection. The implication is that it will be impossible to implement the consistent viewpoint constraint because there will be no specular highlight or lobe due to the Lambertian behaviour. Furthermore, because rough surfaces are considered,

discontinuities in the surface are more likely to be a feature. This does depend on the level of resolution under which the images of the surfaces are captured. However, in this case it is the high frequency nature of the textures which we aim to model and at such a resolution it is unlikely that the surface can be considered smooth and consistent. This makes the integrability constraint considerably less attractive as a candidate for ambiguity resolution. The only viable constraint is the constant vector constraint proposed by Hayakawa. Although our textures do not have constant albedo, the illumination source utilised in the lab is a single moveable light. Since it illuminates the texture in every image we therefore consider the light intensity value to be constant. Applying this constraint will facilitate the reduction of the ambiguity but it does not resolve it completely. There is a residual ambiguity which takes the form of a scaled orthogonal transformation as noted in the table.

5.4 Conclusion

An uncalibrated photometric stereo technique which utilises the constant light source intensity constraint was found to be a promising method with regard to fulfilling our objective of determining surface representations of images of rough surface textures without knowledge of the illumination. The technique requires further development to allow practical implementation due to the residual ambiguity, a scaled orthogonal transform. The resolution of this will be considered in the next chapter.

Chapter 6

Detail of Uncalibrated Photometric Stereo for Texture Planes

6.1 Introduction

In the previous chapter the possibility of achieving our overall goal of determining surface representations by using only image intensity data was explored. The relevant methods were found to primarily concern uncalibrated photometric stereo techniques which attempt to solve the system of equations without information pertaining to the illumination direction. A specific variant of the uncalibrated techniques was identified to be a suitable candidate for use with rough surface textures. However, it was deemed to require further development with particular regard to solving the inherent ambiguity to the extent that the resulting surface estimate images could be utilised for relighting purposes. This work largely entails adapting the approach to both suit and indeed take advantage of its specific application to texture. This involves identifying alternative sources of information to facilitate reduction of the ambiguity with the proviso that they are practical to implement. Detail of the proposed customisation to accomplish this and associated issues will be given in this chapter. A texture-specific uncalibrated photometric stereo technique is thereby derived for use on a practical basis.

This chapter is organised as follows :

Implementation of the uncalibrated photometric stereo technique is reviewed and then developed for texture in Section 6.2. The first stage of the process which concerns the decomposition of the intensity matrix is presented in detail. We then focus on the resolution of the ambiguity to the extent that the resulting surface representation estimates can be utilised for relighting purposes.
6.2 Method

In the previous chapter we identified Hayakawa's [Hayakawa1994] original uncalibrated photometric stereo technique as a potentially viable means of realising our goal of determining surface representations of texture using only intensity data. The issues involved in implementing the technique on a practical basis will be considered in the following discussion. In particular we are concerned with customising the process for use with images of surface textures. It is convenient to consider the overall process in stages and each is defined in flowchart form for clarity (see Figures 6.1-6.4, 6.7).

6.2.1 Decomposition

Our input data consists of a set of f images of a surface texture illuminated from varied but unknown positions. The m pixels of each image are re-arranged into a column. The resulting columns are combined to form a single matrix of intensity values I which is $m \times f$ in size (See Equation 2.34). The fundamental step in the process involves decomposing the resulting intensity matrix to provide initial estimates of the surface representation matrices. This is done by using singular value decomposition which was previously described in Section 2.53. The resulting orthogonal and diagonal matrices, $U\Sigma V^{T}$, are partitioned such that equivalent but smaller matrices are produced which correspond to the first three singular values:

$$\mathbf{U} = \begin{bmatrix} \mathbf{V}' & \mathbf{W}' \\ \mathbf{J} & \mathbf{J}_{-3} \end{bmatrix} \} p \qquad \qquad \mathbf{\Sigma} = \begin{bmatrix} \mathbf{\Sigma}' & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma}'' \\ \mathbf{J} & \mathbf{J}_{-3} \end{bmatrix} \begin{cases} \mathbf{J} & \mathbf{J}_{-3} \\ \mathbf{J} & \mathbf{J}_{-3} \end{cases} \qquad \qquad \mathbf{V}^{\mathrm{T}} = \begin{bmatrix} \mathbf{V}' \\ \mathbf{V}'' \\ \mathbf{J} \end{cases} \}_{f-3}$$

This has the advantage of discarding the noisy part of the input data. Indeed the product of the reduced matrices is the best approximation to an ideal noiseless image intensity matrix which is of rank three and corresponds to an illuminated Lambertian surface (see Equation 6.1).

$$\mathbf{U}'\boldsymbol{\Sigma}'\mathbf{V}'^{T} = \hat{\mathbf{I}}$$
(6.1)

The decomposition components present in this equation are used to give expressions for the *pseudo* surface and illumination matrices which provide initial estimates for the scaled surface normal and illumination vectors (see Equations 6.2, 6.3). This first stage in the uncalibrated photometric stereo algorithm is illustrated in Figure 6.1.

$$\hat{\mathbf{S}} = \mathbf{U}'(\pm \sqrt{\Sigma'})$$
 $\hat{\mathbf{L}} = (\pm \sqrt{\Sigma'})\mathbf{V}'$ (6.2, 6.3)

The sign ambiguity in these pseudo matrices means that there are two valid solutions. These correspond to either left-handed or right-handed co-ordinate systems. Hayakawa advocates selecting the latter in order to satisfy the convention that the *z*-component of the scaled surface normals is positive. This orientation is determined by extracting three scaled surface normals from the two versions of the pseudo surface matrix and using them to form two equivalent matrices. The resulting 3×3 matrix which has a positive determinant is consistent with the right-handed system and allows the appropriate pseudo surface matrix to be identified. It follows that its complementary pseudo illumination matrix is the corresponding choice since Equation 6.1 must be satisfied.



Figure 6.1 Flowchart of uncalibrated photometric stereo first stage which entails formulation of the intensity matrix & its decomposition into pseudo surface & illumination matrices.

6.2.2 Ambiguity Reduction & Resolution

Hayakawa's general approach has been adopted and adapted in related work published subsequently by a number of other authors as is readily apparent from the literature survey a summary of which is given in Table 5.1. The main thrust of this research actually centres on the secondary part of the process which concerns the reduction or resolution of the inherent ambiguity. As previously intimated, the pseudo matrices output from the decomposition stage are one potential solution for the photometric stereo scheme but further information is required to determine the solution uniquely. We need to determine the ambiguity matrix **A** in Equations 6.4 and 6.5 by implementing the procedure outlined in Figure 6.2. This involves acquiring a sufficient amount of appropriate data *by some means* to enable the ambiguity matrix to be estimated. The pseudo estimates of the surface and illumination matrices are subsequently multiplied by this ambiguity matrix to determine the true estimates:

$$S = \hat{S}A$$
 $L = A^{-1}\hat{L}$ (6.4, 6.5)

Ideally this should be achieved without compromising the uncalibrated nature of the algorithm by using alternative sources of information or employing assumptions which can be considered valid in the circumstances. However, the exact nature of the resulting procedure really depends on the application. For example, we are precluded from implementing any of the techniques which are based on specular reflection because we assume that the majority of our textures are largely diffuse in nature.



Figure 6.2 Flowchart of uncalibrated photometric stereo second stage in which the ambiguity is determined in order to transform the pseudo surface & illumination matrices into actual estimates.

Several means of reducing and resolving the ambiguity are discussed in the following:

Constant Light Intensity Assumption

The fact that we utilise a single light source in the capture of the images of real textures for our database means that the constant intensity assumption proposed by Hayakawa is a reasonable approach for our specific application. The theory relevant to this was presented in the previous chapter although it focused on the equivalent constant albedo assumption. The treatment is equivalent in this case since it is merely a matter of considering the size of illumination vectors rather than surface normals.

In practice the value of the intensity parameter will not be known unless it is measured and since this would entail complicating the overall procedure, we opt instead to set it to unity. As noted in the theoretical development, this introduces an unknown constant scaling factor into the solution. However, this is not considered to be an issue as it will actually be taken into account in later stages of ambiguity reduction. Re-arranging the matrix formulation given in the previous chapter to solve for the elements of C, the following is obtained.

$$\begin{bmatrix} c_1 \\ c_2 \\ c_3 \\ c_4 \\ c_5 \\ c_6 \end{bmatrix} = \begin{bmatrix} \hat{l}_{1,x}^2 & 2\hat{l}_{1,x}\hat{l}_{1,y} & 2\hat{l}_{1,x}\hat{l}_{1,z} & \hat{l}_{1,y}^2 & 2\hat{l}_{1,y}\hat{l}_{1,z} & \hat{l}_{1,z}^2 \\ \hat{l}_{2,x}^2 & 2\hat{l}_{2,x}\hat{l}_{2,y} & 2\hat{l}_{2,x}\hat{l}_{2,z} & \hat{l}_{2,y}^2 & 2\hat{l}_{2,y}\hat{l}_{2,z} & \hat{l}_{2,z}^2 \\ \hat{l}_{3,x}^2 & 2\hat{l}_{3,x}\hat{l}_{3,y} & 2\hat{l}_{3,x}\hat{l}_{3,z} & \hat{l}_{3,y}^2 & 2\hat{l}_{3,y}\hat{l}_{3,z} & \hat{l}_{3,z}^2 \\ \hat{l}_{4,x}^2 & 2\hat{l}_{4,x}\hat{l}_{4,y} & 2\hat{l}_{4,x}\hat{l}_{4,z} & \hat{l}_{4,y}^2 & 2\hat{l}_{4,y}\hat{l}_{4,z} & \hat{l}_{4,z}^2 \\ \hat{l}_{5,x}^2 & 2\hat{l}_{5,x}\hat{l}_{5,y} & 2\hat{l}_{5,x}\hat{l}_{5,z} & \hat{l}_{5,y}^2 & 2\hat{l}_{5,y}\hat{l}_{5,z} & \hat{l}_{5,z}^2 \\ \hat{l}_{6,x}^2 & 2\hat{l}_{6,x}\hat{l}_{6,y} & 2\hat{l}_{6,x}\hat{l}_{6,z} & \hat{l}_{6,y}^2 & 2\hat{l}_{6,y}\hat{l}_{6,z} & \hat{l}_{6,z}^2 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}$$

$$\tag{6.6}$$

1

Having determined the elements of the matrix C, it can then be decomposed to provide an estimate of the ambiguity matrix A. Singular value decomposition is utilised for this purpose once again. It is noted that in this case the resulting orthogonal matrices are the same because C is symmetric.

$$\mathbf{C} = \mathbf{W} \boldsymbol{\Pi} \mathbf{W}^{\mathrm{T}} \qquad \mathbf{A} = \mathbf{W} \sqrt{\boldsymbol{\Pi}} \qquad (6.7, 6.8)$$

The initial estimates of the pseudo surface and illumination matrices are multiplied by the resulting ambiguity matrix and its inverse, respectively thus providing secondary estimates (see Figure 6.3). Again, these form a valid solution for the photometric stereo scheme but the solution is still not unique. This is because the number of degrees of freedom have only been reduced (from 9 to 3). A residual ambiguity still exists and further information is required to fully resolve it.



Figure 6.3 Flowchart of initial ambiguity reduction in texture-specific photometric stereo.

Absolute Orientation

The remaining three degrees of freedom amount to an orthogonal transformation as was deduced analytically in the previous chapter. The implication is that the derived pseudo matrices containing estimates of the surface normals and illumination vectors are in an arbitrary co-ordinate system and require alignment to the true co-ordinate system. Horn [Horn1987] presents a number of methods for solving this absolute orientation problem each of which necessitates the knowledge of a number of data points in both co-ordinate systems. Such an approach is neither ideal nor practical for an uncalibrated algorithm since it is likely that this data would have to be obtained by measurement. However, an implementation of this step provides a valuable insight into what is involved with regard to complete resolution of the ambiguity. The least expensive in terms of this requirement for further information will be considered in this section. It requires three data points to be known in each system. Since we are actually concerned with vectors it was initially ventured that one of the three points could be assigned to the origin thus effectively reducing the requisite data burden. However, this was found to be counterproductive with regard to accuracy. We therefore consider the end point of three vectors in this procedure.

For convenience we consider illumination direction in the following and assume that l_1 , $l_2 \& l_3$ have been measured or estimated by some means; we briefly consider ways of achieving this without resorting to calibration later in this section. The vector triad is manipulated to form a local co-ordinate frame in this, the viewer-oriented co-ordinate system which is denoted by the suffix v. The vectors defining the resulting axes of the frame are utilised to form a matrix by arranging them columnwise as shown.

96

$$\mathbf{\dot{x}}_{v} = \mathbf{l}_{2} - \mathbf{l}_{1} \qquad \qquad \mathbf{x}_{v} = \frac{\mathbf{\dot{x}}_{v}}{\|\mathbf{\dot{x}}_{v}\|}$$
(6.9)

$$\mathbf{\dot{y}}_{\nu} = (\mathbf{l}_{3} - \mathbf{l}_{1}) - [(\mathbf{l}_{3} - \mathbf{l}_{1}) \cdot \mathbf{x}_{\nu}]\mathbf{x}_{\nu} \qquad \mathbf{y}_{\nu} = \frac{\mathbf{\dot{y}}_{\nu}}{\|\mathbf{\dot{y}}_{\nu}\|}$$
(6.10)

$$\mathbf{z}_{v} = \mathbf{y}_{v} \times \mathbf{x}_{v} \tag{6.11}$$

$$\mathbf{M}_{v} = \left| \mathbf{x}_{v} \mathbf{y}_{v} \mathbf{z}_{v} \right| \tag{6.12}$$

The corresponding three vectors taken from the derived pseudo matrices $\hat{\mathbf{l}}_1$, $\hat{\mathbf{l}}_2$ & $\hat{\mathbf{l}}_3$ are manipulated in an identical way to generate a matrix \mathbf{M}_a where the suffix *a* denotes the arbitrary co-ordinate system.

$$\mathbf{M}_{a} = \left| \mathbf{x}_{a} \mathbf{y}_{a} \mathbf{z}_{a} \right| \tag{6.13}$$

If $\hat{\mathbf{l}}_{a}$ is a vector from the pseudo illumination matrix, its components along the axes of the local co-ordinate frame are $\mathbf{M}_{a}^{T}\hat{\mathbf{l}}_{a}$. These are mapped into the viewer-oriented co-ordinate system by multiplying by \mathbf{M}_{v} .

$$\mathbf{l}_{v} = \mathbf{M}_{v} \mathbf{M}_{a}^{\mathrm{T}} \hat{\mathbf{l}}_{a}$$
(6.14)

The rotation matrix **R** which resolves the residual ambiguity is therefore written as follows: $\mathbf{R} = \mathbf{M}_{v} \mathbf{M}_{a}^{T}$ (6.15)

This constitutes an orthogonal transformation since each matrix is constructed from the co-ordinate frame axes which are themselves orthonormal. The pseudo surface and illumination matrices are multiplied by \mathbf{R} and its inverse, respectively to generate a unique solution (see Figure 6.4).



Figure 6.4 Flowchart of ideal residual ambiguity resolution in texture-specific photometric stereo.

The significant drawback of this approach to ambiguity resolution is that several viewer-oriented vectors must be known. In terms used to define them, we are obliged to determine three tilt angles and three slant angles. To do so with sufficient accuracy is likely to involve some form of calibration step. Unfortunately this would mean foregoing the very advantage of this photometric stereo technique i.e. its uncalibrated nature. It is possible to estimate the illumination tilt angle in the image by performing a Fourier analysis of the polar plot as proposed by Chantler [Chantler1997] but this is limited to isotropic textures and cannot be applied generally. In any case Horn [Horn1987] highlights a deficiency in this cheap absolute orientation method : it uses the information from the three points selectively such that a different rotation matrix will be obtained if the points (or vectors in this case) are re-ordered. This is not conducive to attaining an accurate solution. Instead Horn recommends tackling the problem in a least squares sense but this would involve knowing an even greater number of viewer-oriented vectors. In summary, it would seem that whilst this method provides a means of resolving the ambiguity, it is not really of practical use in this case due to the data requirement which is unrealistic if the algorithm is to remain a truly uncalibrated one. However, it could still be utilised to test the viability of uncalibrated photometric stereo under ideal conditions. In Chapters 7 & 8 we use this technique (referred to as UPS) as a benchmark for uncalibrated photometric stereo.

Stepwise Orientation

A more practical approach to resolving the residual ambiguity was devised by addressing each of the remaining degrees of freedom individually. In this treatment we consider the overall orthogonal transformation required in three separate stages. This is advantageous since it allows the exact information required for each step to be clearly identified.

z-Axis Alignment

Although the textures described in Chapter 3 have rough surfaces their megastructure is planar. The implication is that given a large enough sample, the mean surface normal of the texture will point directly towards the camera used to capture its image. In terms of the viewer-oriented co-ordinate system we previously defined this means that the mean surface vector points along the *z*-axis and is hence written as $[0,0,1]^T$. This serves to provide further information which can be utilised for orientation at no actual expense in terms of calibration requirement.

Following the ambiguity reduction achieved through the use of the constant intensity assumption, the mean pseudo surface normal is calculated as follows.

$$\overline{\mathbf{n}}_{\text{pseudo}} = \frac{1}{m} \sum_{i=1}^{m} \frac{\hat{\mathbf{s}}_{2(i)}}{\left| \hat{\mathbf{s}}_{2(i)} \right|} = \left[\overline{n}_{\text{pseudo},x} \quad \overline{n}_{\text{pseudo},y} \quad \overline{n}_{\text{pseudo},z} \right]^{\mathrm{T}}$$
(6.16)

$$\overline{\mathbf{n}}_{\text{ideal}} = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^{\mathbf{r}}$$
(6.17)

Taking the dot product of the two vectors enables the angle between them to be determined. This is a measure of the angle of rotation required to orient the pseudo surface vectors such that their mean, a representation of the *z*-axis of the arbitrary co-ordinate system, is aligned with the viewer-oriented *z*-axis. The cross-product yields the corresponding axis of rotation required to perform this operation.

$$\alpha_{1} = \cos^{-1} \left(\overline{\mathbf{n}}_{\text{ideal}} \cdot \overline{\mathbf{n}}_{\text{data}} \right)$$
(6.18)

$$\mathbf{\dot{w}}_{1} = \overline{\mathbf{n}}_{\text{ideal}} \times \overline{\mathbf{n}}_{\text{data}}$$
 $\mathbf{w}_{1} = \frac{\mathbf{\dot{w}}_{1}}{|\mathbf{\dot{w}}_{1}|}$
(6.19)

Inserting these into the general formula for a rotation matrix [Woo1999] allows the required transform to be determined.

$$\mathbf{R}_{1} = \mathbf{w}_{1}\mathbf{w}_{1}^{\mathrm{T}} + \cos(\alpha_{1})(\mathbf{I}_{d} - \mathbf{w}_{1}\mathbf{w}_{1}^{\mathrm{T}}) + \sin(\alpha_{1})\begin{bmatrix} 0 & -w_{1z} & w_{1y} \\ w_{1z} & 0 & -w_{1x} \\ -w_{1y} & w_{1x} & 0 \end{bmatrix}$$

where $\mathbf{I}_{d} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ (6.20)

The resulting rotation matrix is used to update the estimates for both pseudo surface normal and pseudo illumination matrices accordingly.

$$\hat{\mathbf{S}}_{3} = \hat{\mathbf{S}}_{2}\mathbf{R}_{1}$$
 $\hat{\mathbf{L}}_{3} = \mathbf{R}_{1}^{-1}\hat{\mathbf{L}}_{2}$ (6.21, 6.22)

The effect of this operation is readily perceived from the surface normal distribution graphs presented in Figure 6.5.

Having aligned the *z*-axis of the arbitrary co-ordinate system with that of the vieweroriented system, there remain two degrees of freedom regarding the residual ambiguity. One of these concerns the orientation of the *x*-*y* plane and will be considered in the following subsection.



Figure 6.5 Surface normal distribution for synthetic fractal texture corresponding to pseudo matrix estimate (a) before & (b) after z-axis alignment.

Rotation about *z***-Axis**

With regard to the orientation of the x-y plane, the objective is to transform the pseudo vectors such that the x-axis of their local co-ordinate system (and therefore the y-axis too) corresponds to that of the viewer-oriented system without perturbing the mean surface vector direction which has been aligned. This is achieved by performing a rotation about the z-axis itself. Since the axis of rotation is known, it is matter of determining the angle of rotation to generate the required matrix. We note that an image registration step would not be applicable with regard to finding a solution in this case. This is because the rotation concerns the individual vectors and normals rather than their spatial positioning. The tilt angle corresponding to a single viewer-oriented vector must be known. If the texture is isotropic then Chantler's method [Chantler1997] can be utilised to estimate a value for the tilt angle of an illumination vector from one of the input images. However, in the more general case it will be necessary to supply a tilt angle which has been measured or is known. The unfortunate consequence of this is that although the information burden is extremely low, it does mean that we are forced to renege on our ideal of a truly uncalibrated photometric stereo algorithm. However, we believe that setting up the equipment such that the initial placement of the light corresponds to a tilt angle of 0° is not unreasonable and could easily be implemented on a practical basis. We feel that this approach is preferable in order to retain the general applicability of the algorithm rather than limit it to specific textures such as those with smooth surfaces or exhibiting specular reflection. It may actually be sufficient to assign one image as a reference and hence designate its corresponding illumination tilt angle as 0° . The implication is that the illumination direction corresponding to a relit image would be relative to the allocated co-ordinate system. This may be acceptable for some applications.



Figure 6.6 Pseudo illumination vector corresponding to known vector requires rotation of α_2 to orient x-y plane with viewer-oriented co-ordinate system.

The tilt angle corresponding to the first vector in the pseudo illumination matrix is given by the following equation:

$$\tau_{pseudo} = \cos^{-1} \left(\frac{\hat{l}_{3(1),x}}{\sin(\cos^{-1}(\hat{l}_{3(1),z}))} \right)$$
(6.23)

Since the corresponding known tilt is 0° , the required angle of rotation is straightforward to deduce (see Figure 6.6).

$$\alpha_2 = \tau_{known} - \tau_{pseudo} = -\tau_{pseudo} \tag{6.24}$$

This angle is substituted into the generic *z*-axis rotation formula [Woo1999] to give the required rotation matrix. It is noted that the inverse is specified here since we are considering illumination vectors and our convention has been to specify the ambiguity matrix corresponding to the pseudo surface normal matrix as \mathbf{R} .

$$\mathbf{R}_{2}^{-1} = \begin{bmatrix} \cos(\alpha_{2}) & -\sin(\alpha_{2}) & 0\\ \sin(\alpha_{2}) & \cos(\alpha_{2}) & 0\\ 0 & 0 & 1 \end{bmatrix}$$
(6.25)

The resulting rotation matrix is used to update the estimates for both pseudo surface normal and pseudo illumination matrices accordingly.

$$\hat{\mathbf{S}}_4 = \hat{\mathbf{S}}_3 \mathbf{R}_2$$
 $\hat{\mathbf{L}}_4 = \mathbf{R}_2^{-1} \hat{\mathbf{L}}_3$ (6.26, 6.27)

The co-ordinate system corresponding to the resulting pseudo matrix estimates is now aligned with that of the viewer-oriented system. The remaining degree of freedom concerns the overall scaling of the vector field. In practical terms further processing could therefore be considered to be unnecessary with regard to the focus of this work, visualisation applications. This is because the designer utilising the resulting bump map to boost photorealism is likely to dictate the exact scale required for the 3D virtual scene. With regard to the quantitative assessment of the technique, however, we are obliged to consider this issue in order to allow comparison with the ground-truth data set i.e. the input images. This will be considered in the next subsection.

Overall Scaling

In order to resolve the ambiguity completely and hence facilitate quantitative assessment, it is necessary to transform the co-ordinate system of the pseudo estimates whose axes are now aligned with the viewer-oriented system such that the axes are of equivalent magnitude. The scaling factors involved must therefore be determined to resolve this final degree of freedom. To do so, a single slant angle corresponding to a vector in the viewer-oriented system must be known. For convenience, we assume that once again this corresponds to the first illumination vector.

The slant angle corresponding to the first vector in the pseudo illumination matrix is given by the following equation.

$$\sigma_{pseudo} = \cos^{-1}(\hat{l}_{4(1),z})$$
(6.28)

The resulting scaling factors corresponding to each axis are therefore given by the following ratios.

$$factor_{x} = factor_{y} = \frac{\sin(\sigma_{known})}{\sin(\sigma_{pseudo})}$$
(6.29)

$$factor_{z} = \frac{\cos(\sigma_{known})}{\cos(\sigma_{pseudo})}$$
(6.30)

These factors are used to form the required transformations which take the form of diagonal matrices.

$$\mathbf{R}_{3}^{-1} = \begin{bmatrix} factor_{x} & 0 & 0\\ 0 & factor_{y} & 0\\ 0 & 0 & factor_{z} \end{bmatrix}$$
(6.31)

$$\mathbf{R}_{3} = \begin{bmatrix} factor_{x}^{-1} & 0 & 0\\ 0 & factor_{y}^{-1} & 0\\ 0 & 0 & factor_{z}^{-1} \end{bmatrix}$$
(6.32)

The actual estimates of the surface normal matrix and the illumination matrix can finally be determined.

$$S = \hat{S}_4 R_3$$
 $L = R_3^{-1} \hat{L}_4$ (6.33, 6.34)

A flowchart of the proposed process to resolve the residual ambiguity is given in Figure 6.7.



Figure 6.7 Flowchart of practical residual ambiguity resolution in texture-specific photometric stereo.

The scaled surface normal matrix S is used to produce surface gradient and albedo images which can be utilised for relighting such that colour images of the texture are generated under user-specified illumination conditions. The resulting images can hence be compared with the original images of the texture captured by digital camera to facilitate a quantitative assessment of the proposed algorithm.

6.3 Summary & Discussion

In this chapter we have developed an algorithm which is based on Hayakawa's original method of uncalibrated photometric stereo for specific use with the images of illuminated rough surface textures. The main issue was noted to be the inherent ambiguity which corresponds to nine degrees of freedom. Appropriate methods were presented to reduce and subsequently resolve these in order to work towards obtaining a unique solution for both surface and illumination matrices. The associated issues with the various approaches were discussed. Initially we proposed to reduce the ambiguity from nine to three degrees of freedom through the use of the constant light source assumption advocated by Hayakawa. This was deemed to be reasonable because a single moveable light source was utilised in the capture of the database of images of illuminated real textures. In order to resolve the remaining three degrees of freedom we initially considered Horn's simplest absolute orientation method. Whilst this allows our objective of ambiguity resolution to be achieved, the information burden in terms of viewer-oriented vectors was deemed to be too high to justify its use in an uncalibrated algorithm. Instead, the fact that the remaining degrees of freedom constitute an orthogonal transformation led us to consider the orientation in separate rotation and scaling steps. The first stage involves orienting the z-axis of the arbitrary co-ordinate system estimated by the mean pseudo surface vector with that of the viewer-oriented system. The next stage entails the orientation of the x-y plane. This necessitates the knowledge of a single tilt angle of an actual illumination or surface vector. With regard to the former, the use of Chantler's tilt estimation technique was suggested as a way of determining a value for this but was noted to be applicable to isotropic textures only. In the more general case e.g. when textures are directional, we advocated setting up the image capture equipment such that the first illumination vector tilt angle is known. We acknowledged the fact that this forces us to relinquish the uncalibrated status of the algorithm. However, in practice it may well be reasonable to simply assign a tilt angle of 0° to the first image and hence avoid the need for calibration. In doing so the resulting surface representation would be relative to the chosen co-ordinate system but this may be acceptable for relighting applications. The third stage is a scaling step. Whilst this requires the value of a slant angle corresponding to an actual illumination or surface vector to be known, it was noted that it is only necessary to perform this operation for quantitative assessment of the algorithm.

Overall, we have described a plausible implementation of an algorithm which is based on Hayakawa's uncalibrated photometric stereo. This texture-specific method simply requires the input of intensity data corresponding to a series of images of a texture illuminated from different directions with a common light source and a single illumination vector in order to generate estimates of the scaled surface normal and illumination matrices. We feel that the approach we have adopted for ambiguity reduction/resolution is preferable to the alternative of making further assumptions which restrict the general applicability of the algorithm.

Chapter 7

Assessment of Uncalibrated Photometric Stereo with Texture – Simulation

7.1 Introduction

In the previous chapter a scheme based on uncalibrated photometric stereo was developed for use with surface texture. In this chapter the proposed technique is examined with regard to performance under various simulated conditions. We gauge the robustness of the algorithm for non-ideal reflectance and in the presence of noise. The effect of both number and position of illuminants is also evaluated with a view to recommending operating conditions which provide high accuracy. An identical analysis is carried out for variants of the photometric stereo technique to allow comparison and hence provide an assessment of relative performance.

This chapter is organised as follows:

Variants of the photometric stereo algorithm which were utilised in the simulation investigation are reviewed in Section 7.2. Experimental results are then reported in Section 7.3. The findings are summarised and conclusions are drawn in Section 7.5.

7.2 Experimental Technique Variations

To facilitate an assessment of the results for the texture-specific algorithm, we conducted parallel experiments with up to four other variants of the Lambertian photometric stereo algorithm for comparison. Each of the algorithms utilised is described in brief in the following.

PS-3

This is the three-image version of photometric stereo [Woodham1980] which is described in Section 2.5.3. It is a calibrated technique and involves determining the inverse of the illumination matrix.

PS

This is the over-constrained version of photometric stereo [Woodham1980] which is described in Section 2.5.3. It is a calibrated technique and involves determining the pseudo-inverse of the illumination matrix in order to find the least squares solution.

UPS

This is the uncalibrated photometric stereo technique originally proposed by Hayakawa which is described in Section 6.2 [Hayakawa1994]. It uses SVD to decompose the intensity matrix. It reduces the ambiguity by assuming constant light source intensity such that the illumination vectors are constrained to be of constant magnitude. It resolves the residual orthogonal ambiguity by using an absolute orientation method [Horn1987]. In this case the technique strictly loses its uncalibrated status because absolute orientation demands the knowledge of three illumination vectors. However, its use provides us with a baseline against which the texture-specific algorithm may be compared.

UPS-tx

This is the texture-specific uncalibrated technique which we developed in the previous chapter. It is described in Section 6.2. Similar to the UPS method, it uses SVD to decompose the intensity matrix and it reduces the ambiguity by assuming constant light source intensity. It resolves the residual ambiguity by using a stepwise approach rather than absolute orientation, however. First of all, the mean surface normal of the pseudo matrix is aligned with that known from the fact that the textures are constrained to have a planar megastructure. The tilt and slant angle of one illumination direction must be known to facilitate a further rotation and scaling of the pseudo matrix vectors.

UPS-it

In this case we modified the UPS-tx method by supplementing it with a pre-processing step in which outlying intensity data is discarded. This is based on a method by Georghiades which attempts to address issues such as the presence of shadows and highlights in the intensity matrix [Georghiades2003a]. This was introduced in Section 2.5.3. In this case we focus on shadowing and set a lower bound for intensity value for the iterative procedure.

Despite their requirement for a number of illumination vectors to be known, we feel that it is reasonable to label UPS, UPS-tx and UPS-it techniques as 'uncalibrated'. A term such as 'pseudo-uncalibrated' or 'semi-uncalibrated' may be more accurate in this case. However, the general approach of intensity matrix decomposition followed by ambiguity reduction and resolution is usually referred to as 'uncalibrated' in the literature [Drbohlav2002, Drbohlav2003, Georghiades2003c].

7.3 Technique Performance

The Lambertian photometric stereo algorithms (PS-3, PS, UPS, UPS-tx & UPS-it) are based on an assumption of perfectly diffuse reflection from a surface illuminated by a point light source at infinity. We decided to explore the behaviour of the proposed technique (UPS-tx) under reflectance conditions which challenge this assumption and compare its performance with those of the other algorithms. Such conditions include the presence of shadows and specular highlights in addition to the case when the point light source is nearby. These issues are considered in Sections 7.3.1, 7.3.2 and 7.3.3 respectively. We investigate the effect of the presence of noise (Section 7.3.4) because this will be an inherent feature in the images of real textures captured by camera. In Chapter 4 we determined that the configuration of the illumination vectors affects the performance of the three-image photometric stereo algorithm (PS-3). We examine the effect of illuminant number and relative position on the proposed uncalibrated technique in Section 7.3.5. Finally we evaluate the sensitivity of the algorithm with regard to its mean surface vector assumption. With regard to all of the aforementioned experiments, we use the same data set as input for the five algorithms in each case to facilitate an accurate comparison of the uncalibrated and calibrated techniques. A randomly illuminated series of texture images would be the likely data input for the uncalibrated techniques in practice. However, we decided that this was not appropriate for simulation considering our objective of gauging relative performance accurately.

7.3.1 Shadowing

Motivation

Our proposed technique (UPS-tx) is based on ideal diffuse reflection and hence assumes that every pixel in each of the input images is illuminated. In reality this will not be the case for many real textures whose mesostructure is three-dimensional. Shadows occur when the incoming ray of light is occluded and this is more likely for surfaces with greater variation in height. It is important to establish the performance of the algorithm under these conditions since shadows will be encountered in practice.



Figure 7.1 Visual explanation of both self and cast shadowing.

Data Generation

For simulation purposes, a means of modelling shadows is required. Two different processes are commonly considered to account for the occlusion of the incident light beam which results in the presence of shadows. These are cast shadowing and self shadowing [Schlüns1997]. The former refers to surface facets which lie in the shadow projected by a relatively high neighbouring facet which obscures the light source. The latter refers to the situation when the angle between the surface facet normal and the illumination direction is greater than 90° i.e. when the light source is below the facet horizon. See Figure 7.1 for an illustration. Modelling self shadowing merely involves setting the intensity value to zero when the value of the dot product between the surface facet normal and the illumination vector is negative:

$$i(x,y) = \begin{cases} \rho(x,y)\lambda \mathbf{n}(x,y) \cdot \mathbf{b} & 0 \le \theta \le 90^{\circ} \\ 0 & \theta > 90^{\circ} \end{cases}$$
(7.1)

Cast shadowing is more demanding to implement and has been ignored for the purposes of this investigation. In this section we present results which are based solely on the presence of self shadows.



Figure 7.2 Sample height map profile for fractal surface.

For our experiments we created a series of data sets for each of the three synthetic surface textures detailed in Section 3.4. Although self shadowing depends on the relative orientation of the surface normal and the illumination vector, we opted for consistent illumination directions as specified in Section 3.4.2 and controlled the extent of shadowing by altering the surface roughness. This was achieved by multiplying the height map corresponding to each synthetic surface by a factor greater than unity and differentiating the resulting image to obtain the new surface gradient maps of an equivalent but rougher surface (see Figure 7.2). The degree of surface roughness was estimated for each surface by calculating the rms roughness z_{rms} parameter which is defined in Section 3.3. Five values of rms roughness ($0 < z_{rms} < 20$) were employed for each synthetic texture.

Equation 7.1 was used to generate images of each surface under the range of illumination directions corresponding to: $0^{\circ} \le \tau < 360^{\circ}$, $\Delta \tau = 10^{\circ}$, $30^{\circ} \le \sigma \le 60^{\circ}$, $\Delta \sigma = 15^{\circ}$. Each data set therefore contained 108 images. See Figure 7.3 for examples. These images provide a visual demonstration that an increase in height variation results in a higher occurrence of shadows since the rougher textures are relatively dark. Equivalent data sets were also generated with Equation 2.2 for comparison. In this case negative intensity values are permitted.



Figure 7.3 Sand ripple surface of increasing surface roughness illuminated under identical conditions. Image (c) is the roughest.

Method

Thirty-six images corresponding to a constant illumination slant angle σ of 45° were selected from each data set for use as input data. These were processed with the five Lambertian photometric stereo techniques described in Section 7.2. The resulting *p* and *q* maps and albedo image were used in conjunction with the Lambertian model to produce seventy-two relit images with illumination directions not present in the input data i.e. $0^{\circ} \le \tau < 360^{\circ}$, $\Delta \tau = 10^{\circ}$, $\sigma = 30^{\circ}/60^{\circ}$. A signal to relight error ratio (*SER*) value was calculated for each relit image using Equation 3.9. A texture signal to relight error ratio (*TSER*) was subsequently calculated for each data set of each synthetic texture using Equation 3.10. Finally, a mean *TSER* value for the three synthetic textures was calculated for each data set.

Results

With regard to the UPS-tx algorithm, Figure 7.4 demonstrates that when negative intensity values are permitted the performance is comparatively unaffected by surface roughness. This is not surprising because the data is effectively perfect. If self shadows are modelled, however, the accuracy in terms of the mean *TSER* value decreases as the rms roughness increases. This can be attributed to the increasing presence of shadows which are a source of error for the algorithm. Their effect can be observed from a visual inspection of the albedo image which should be of constant intensity (see Figure 7.5).

A comparison of the relight accuracy for the five techniques described in Section 7.2 is presented in graphical form in Figure 7.6. The fact that the plots have a similar profile to each other implies that the increasing presence of shadows has a negative impact on the performance of Lambertian photometric stereo in general. We observe that the UPS-tx algorithm is practically as robust to surface roughness as the technique with the best performance, PS.



Figure 7.4 Mean TSER for three synthetic textures versus rms roughness for the UPS-tx algorithm.



Figure 7.5 Sample output images for sand ripple surface (rms roughness=5).



Figure 7.6 Mean TSER for three synthetic textures versus rms roughness for five Lambertian photometric stereo algorithms.

7.3.2 Specularities

Motivation

When reflection from a surface is specular in nature, highlights are observed. These correspond to large spikes in image intensity value. The Lambertian model is unable to accurately represent specular reflection. This is illustrated in Figure 7.7. As a result, textures which exhibit specular reflectance are not ideal candidates for use with Lambertian photometric stereo techniques. However, it is important to establish the performance of the algorithm when presented with images of such textures since they may well be encountered in practice.



Figure 7.7 Reflection from a surface facet (a) diffuse reflection modelled by Lambert's law, (b) specular reflection modelled by the Phong model with an exponent of 10 and (c) specular reflection modelled by the Phong model with an exponent of 100.

Data Generation

For our experiments we created a series of data sets for each of the three synthetic surfaces detailed in Section 3.4. The simplest reflection model describing specular reflectance is the equation proposed by Phong [Phong1975]. It was introduced in Chapter 1 and is given by Equation 1.3. The equation was used to generate images of each surface under the range of illumination directions corresponding to: $0^{\circ} \le \tau < 360^{\circ}$, $\Delta \tau = 10^{\circ}$, $30^{\circ} \le \sigma \le 60^{\circ}$, $\Delta \sigma = 15^{\circ}$. The data sets encompass a range of values for specular reflection proportion (0-100%) and also a range of the Phong exponent (0-100).

Method

Thirty-six images corresponding to a constant illumination slant angle σ of 45° were selected from each data set for use as input data. These were processed with the five Lambertian photometric stereo techniques described in Section 7.2. The resulting *p* and *q* maps and albedo image were used in conjunction with Equation 7.1 to produce seventy-two relit images with illumination directions not present in the input data i.e. $0^{\circ} \le \tau < 360^{\circ}$, $\Delta \tau = 10^{\circ}$, $\sigma = 30^{\circ}/60^{\circ}$. A signal to relight error ratio (*SER*) value was calculated for each relit image using Equation 3.9. A texture signal to relight error ratio (*TSER*) was subsequently calculated for each data set of each synthetic texture using Equation 3.10. Finally, a mean *TSER* value for the three synthetic textures was calculated for each data set.

Results

With regard to the UPS-tx algorithm, the plot in Figure 7.8 shows that in general its performance deteriorates as the proportion of specular reflection increases. However, it is also apparent that the algorithm is relatively robust for low levels of specular reflection. The effect of the Phong exponent on performance is more readily discerned by re-plotting the data (see Figure 7.9).



Figure 7.8 Mean TSER for three synthetic textures versus proportion of specular reflection over a range of Phong exponent values for the UPS-tx algorithm.



Figure 7.9 Mean TSER for three synthetic textures versus Phong exponent over a range of specular reflection proportions for the UPS-tx algorithm.

For lower proportions of specular reflection (< 0.6) we observe that as the value of the exponent increases the accuracy initially falls but then rebounds to even higher levels. This can be explained by considering the exposure of the viewer in the scene to the spike in intensity value which is characteristic of specular reflection (see Figure 7.10).



Figure 7.10 Phong reflection with exponent of (a) 2, (b) 10 and (c) 100. Graphs generated from spreadsheet created by Dr. M. J. Chantler.

For high values of the exponent, the specular component will correspond to a narrow intensity peak (see Figure 7.10c). This means that the highlight will only be observed by the viewer for an extremely limited range of illumination angles. In this case diffuse reflection will be dominant. For low values of the exponent the peak is extremely broad (see Figure 7.10a). The fact that it will be observed over a large range of illumination angles actually serves to dilute the effect. The specular component tends towards sinusoidal behaviour as the exponent is reduced to unity. In this case the overall



Figure 7.11 Scaled images of a Mulvaney surface exhibiting (a) diffuse reflection, (b) specular reflection (Phong exponent = 10) and (c) specular reflection (Phong exponent = 100).

reflection behaviour approximates to that of the Lambertian model although this corresponds to a pseudo surface and not the true surface. Intermediate values of the exponent result in a peak of medium breadth which may be observed over a reasonable range of illumination angles (see Figure 7.10b). In this case the effect of specular reflection is much more likely to be discerned by the viewer. This behaviour can also be observed from a visual comparison of illuminated surface images (see Figure 7.11). It is evident that the central image (Figure 7.11b) which corresponds to an intermediate exponent value is relatively brighter than the two other images. This signifies the presence of more visible highlights. The image corresponding to the highest exponent value (Figure 7.11c) resembles the image of the diffuse surface (Figure 7.11a). This would clearly be beneficial with regard to the performance of the UPS-tx algorithm since it assumes Lambertian reflectance.

Examples of output images for the data set corresponding to Figure 7.11b are given in Figure 7.12. Since the albedo used in the generation of the data set images is constant the fact that its estimate is not uniform emphasises that the presence of highlights provides a source of inaccuracy.



Figure 7.12 Sample output images for specularly reflecting Mulvaney surface data set corresponding to a Phong exponent of 10 and a specular reflection proportion of 50%.

With regard to comparative performance of the five techniques, the PS algorithm attains the best performance for specular reflectance relative to the UPS-tx, UPS and PS-3 algorithms (see Figure 7.13 and Figure 7.14). It is evident that the UPS-tx algorithm is almost as robust to these conditions since its performance profile is similar to that for the PS algorithm. The UPS-it algorithm was tested on the same range of data sets for these specular reflection experiments. However, in many instances convergence of the iteration was problematic and as such it became impractical to implement. We found that input data which diverges significantly from diffuse reflection is unsuitable for this algorithm. This is not surprising because the algorithm has been designed to provide the best-fit solution to a Lambertian model.



Figure 7.13 Mean TSER for three synthetic textures versus proportion of specular reflection for four Lambertian photometric stereo techniques (Phong exponent=40).



Figure 7.14 Mean TSER for three synthetic textures versus Phong exponent for four Lambertian photometric stereo techniques (specular reflection proportion=0.4).

7.3.3 Point Illumination

Motivation

The Lambertian photometric stereo techniques assume that the surface is lit by a point light source at infinity. The implication is that the resulting rays of light are parallel with each other when they impinge on the surface and it is hence possible to treat the illumination vector as constant. The distance from the light source to the surface is finite in practice, however. In this case the illumination vector varies over the surface (see Figure 7.15). The degree to which this happens depends on the physical dimensions of the experimental set-up. It is important to investigate this to determine the effect of nearby point illumination on the performance of the UPS-tx algorithm.



Figure 7.15 Point illumination diagram illustrating difference in illumination vectors with regard to both magnitude and direction.

Data Generation

Simulating nearby point light source illumination involves using a version of the Lambertian model which is explicit in terms of the illumination radius r_b (see Equation 7.2). Image intensity is inversely proportional to the square of this distance; when the source is at infinity r_b can be taken as constant and is incorporated into the albedo term.

$$i(x,y) = \frac{\rho(x,y)\lambda}{r_b^2(x,y)} \mathbf{n}(x,y) \cdot \mathbf{b}(x,y)$$
(7.2)

The intensity also depends on the dot product of the surface normal and the illumination vector. Since the direction and magnitude of the illumination vector vary over the surface in this case, they must be computed for each facet in order to implement point illumination (see Figure 7.16). This requires the position of the light source to be known relative to the surface. To establish this, we specified the dimensions of the surface i.e. physical distance and in number of pixels. We also specified the length of the light beam impinging on the centre of the surface which corresponds to the principal illumination direction (see Figure 7.16a). This approach was used to generate images of

each synthetic surface under point illumination for the range of illumination directions specified in Section 3.4.2. In this case we consider these directions to refer to the principal illumination vector. The data sets encompass a range of illumination radii and surface areas ($r_b(x_{P},y_{P}) = 0.5$ m-2.0m; $s_{tx} = 0.1$ m & 0.4m).



Figure 7.16 Point illumination calculations (a) light source co-ordinate, (b) illumination angles and radius for a facet.

Method

Thirty-six images corresponding to a constant illumination slant angle σ of 45° were selected from each data set for use as input data. These were processed with the five Lambertian photometric stereo techniques described in Section 7.2. The resulting *p* and *q* maps and albedo image were used in conjunction with Equation 7.1 to produce seventy-two relit images with illumination directions not present in the input data i.e. $0^{\circ} \le \tau < 360^{\circ}$, $\Delta \tau = 10^{\circ}$, $\sigma = 30^{\circ}/60^{\circ}$. A mean *TSER* value for the three synthetic textures was calculated for each data set using Equations 3.9-3.11.

Results

With regard to the plots in Figure 7.17, it is clear that performance of the UPS-tx algorithm gradually deteriorates as the illumination radius is shortened. Furthermore, the relight error is relatively higher for physically larger surfaces. These findings can be explained by considering the deviation from the reflectance model which is assumed by the technique. The key factor is the variability in the illumination gradient over the surface of the texture. The unscaled images in Figure 7.18a provide a visual demonstration that a greater illumination gradient occurs when the light source is relatively close to the texture surface. This has a negative impact on performance. With regard to surface size, it is clear that the illumination gradient will be reduced if a smaller sample of the surface is considered. Performance could be improved in this way.

(a) 25 dB - PS-3 - PS 20 dE UPS UPS-ty UPS-it 15 dE Mean TSER 10 dB 5 dB 0 dB 0.5 1 1.5 2 Illumination radius (m) (b) 25 dB 20 dB Mean TSER 15 dB 10 dB ···· PS-3 PS 5 dB UPS UPS-tx UPS-it 0 dB 0.5 1 1.5 2 Illumination radius (m)

Figure 7.17 Mean TSER for three synthetic textures versus illumination radius for square sample of side (a) 0.4m & (b) 0.1m for five photometric stereo algorithms.



Figure 7.18 Images of fractal surfaces illuminated by a nearby point light source with illumination radius between 0.5m and 2.0m (a) images not scaled with each other to facilitate illumination gradient evaluation (b) images scaled with each other.



Figure 7.19 Sample output images for data set corresponding to square fractal surface of side 0.4m under point illumination of radius 0.5m.

Examples of output images for the data set corresponding to the image with a 0.5m illumination radius in Figure 7.18 are given in Figure 7.19. Since the albedo used in the generation of the data set images is constant the fact that its estimate is not uniform emphasises the adverse effect of the illumination gradient. This is also noticeable in the p and q map images.

With regard to comparative performance (see Figure 7.17), it is evident that the UPS-tx algorithm is as robust to the effect of nearby point illumination as the calibrated Lambertian photometric stereo techniques PS and PS-3. It is also apparent that UPS-tx outperforms the standard uncalibrated technique, UPS. This is particularly noticeable in Figure 7.17b. However, this behaviour is not merely a feature of these point illumination experiments since it may also be observed in the previous sections on closer inspection (see Figures 7.6, 7.13, 7.14). We attribute the relatively poor performance of the UPS algorithm to its use of absolute orientation with three data points [Horn1987] for ambiguity resolution. Accuracy concerns regarding this rudimentary technique were highlighted in Chapter 6 and it is therefore likely to be the main factor with regard to the poorer performance. The iterative technique UPS-it was found to have the poorest performance overall. We attribute this to the fact that the input images are increasingly darker as the illumination radius is increased (see Figure 7.18b). This means that a higher percentage of intensity data will be considered as outlying and initially discarded. It is evident that the algorithm is sensitive to a lack of valid intensity values. In practice this problem would be resolved by opening up the aperture of the camera in order to obtain a suitable range of intensities.

7.3.4 Presence of Noise

Motivation

Noise is an inherent feature in the images of real textures. It can be generated by a number of mechanisms which take place during the imaging process [McGunnigle1998]. The use of a CCD camera to produce images means that sensor noise is likely to be a contributory factor in this regard. These cameras employ a semiconductor sensor to measure irradiance; we assume that the intensity values in the resulting images have a direct linear relationship to the irradiance. However, if the light falling on the sensor exceeds certain levels then processes such as signal attenuation, clipping and blooming are likely to occur [Klette1999]. These processes result in a non-linear response. The fact that CCD chips generate charge when no light is falling on them is a further source of error. This means that in the absence of light, an intensity value known as the black level will register in the image. The consequence of sensor noise is that the image intensity values output from the camera deviate from those This has implications for the Lambertian photometric stereo actually observed. algorithms because they assume that the intensity data results from ideal diffuse reflection. It is important to determine how robust the UPS-tx algorithm is to noisy input data.

Data Generation

McGunnigle determined that the temporal noise associated with the images of textures could be approximated by a white noise process [McGunnigle1998]. Noise can be modelled by using a power spectral density which is inversely proportional to the frequency raised to a power β . This is analogous to the fractal given by Equation 3.5 in Section 3.4.1. For white noise the roll-off factor β is set to zero. We generated a series of white noise images which are Gaussian with zero mean. These were added to the images generated by relighting each synthetic surface using Equation 7.1 under the range of illumination directions specified in Section 3.4.2. The resulting data sets encompass a range of input signal to noise levels (0 - 25dB). See Figure 7.20 for examples.



Figure 7.20 Images of (a) white noise, (b) illuminated fractal surface and (c) image resulting from the addition of (a) and (b). Signal to noise ratio is ~ 10 dB.

Method

Thirty-six images corresponding to a constant illumination slant angle σ of 45° were selected from each data set for use as input data. These were processed with four Lambertian photometric stereo techniques (PS-3, PS, UPS, UPS-tx); we decided to discontinue the use of the UPS-it technique at this juncture due to its record of poor performance. The resulting *p* and *q* maps and albedo image were used in conjunction with Equation 7.1 to produce seventy-two relit images with illumination directions not present in the input data i.e. $0^{\circ} \le \tau < 360^{\circ}$, $\Delta \tau = 10^{\circ}$, $\sigma = 30^{\circ}/60^{\circ}$. *TSER* values were calculated for each data set of each synthetic texture using Equations 3.9-3.10.

Results

In this section we present results for the fractal surface only. The general profile is similar for the three surfaces, however. With regard to Figure 7.21, it is clear that increasing levels of noise in the input data have a detrimental effect on the performance of Lambertian photometric stereo in general. The UPS-tx algorithm attains comparable levels of accuracy to the other methods although the PS algorithm is superior. Examples of output images for a noisy fractal data set are given in Figure 7.22. Since the albedo used in the generation of the images is constant the fact that the estimated albedo image is not uniform emphasises the adverse impact of noise.



Figure 7.21 TSER versus input signal to noise ratio for an illuminated fractal surface for four Lambertian photometric stereo techniques.



Figure 7.22 Sample output images for the fractal data set corresponding to an input signal to noise ratio of 10dB.

7.3.5 Illuminant Number & Positioning

Motivation

Multiple images of a surface texture each illuminated from a different direction are required as input for the UPS-tx algorithm. Whilst a minimum of six images must be used due to the orientation requirements, there is no upper limit on their number. The exact number of images and therefore illumination directions utilised may influence the performance of the algorithm. The existence of an optimal illumination configuration for the three image technique (PS-3) suggests that this could also be an important factor in this case. Although the UPS-tx algorithm is uncalibrated and illumination directions which are beneficial in terms of performance.

Data Generation

We created a data set for each of the three synthetic surface textures detailed in Section 3.4 by using Equation 7.1 to generate images of each surface under the range of illumination directions corresponding to: $0^{\circ} \le \tau < 360^{\circ}$, $\Delta \tau = 10^{\circ}$, $30^{\circ} \le \sigma \le 60^{\circ}$, $\Delta \sigma = 15^{\circ}$. Each data set therefore contained 108 images.

Method

We adopted two approaches with regard to the experiments in this case. One involved varying the number of input images such that the corresponding illumination configuration was always symmetric. The second involved varying the number of input images such that the configuration became asymmetric. In this case only the PS, UPS and UPS-tx techniques were utilised; the PS-3 technique only uses three input images and therefore cannot be employed for the purpose of these experiments.

With regard to the asymmetric investigation thirty-six images corresponding to a constant illumination slant angle σ of 45° were selected from each data set for use as the *initial* input data. The first application of the three algorithms in this case therefore actually involved data with a symmetric configuration. Subsequent applications of each algorithm involved monotonically reducing the number of input images such that the tilt angle range decreased by 30° each time (0° $\leq \tau \leq 360^{\circ}$ -k30° for k=1-8, $\Delta \tau$ =10°). For the symmetric investigation thirty-six images corresponding to a constant illumination slant angle σ of 45° were again selected from each data set for use as the *initial* input data. Subsequent applications of each algorithm involved monotonically reducing the number of input images such that the illuminant configuration was maintained in a symmetric pattern. In each case the resulting p and q maps and albedo image were used in conjunction with Equation 7.1 to produce seventy-two relit images with illumination directions corresponding to: 0° $\leq \tau < 360^{\circ}$, $\Delta \tau=10^{\circ}$, $\sigma=30^{\circ}/60^{\circ}$. A mean *TSER* value for

the three synthetic textures was calculated for each configuration using Equations 3.9-



Figure 7.23 Mean TSER for three synthetic surfaces versus number of images input to three photometric stereo techniques. Illuminant configurations largely asymmetric.

Results

With regard to the asymmetric investigation, it is evident that using fewer images is detrimental to the performance of the three algorithms tested (see Figure 7.23). The UPS-tx algorithm is the most sensitive in this case. The calibrated technique PS is the most robust to these conditions. Its performance is relatively unaffected by reducing the number of input images from 36 to 28. A very different performance profile is generated when the number of input images is decreased but the corresponding illumination configuration is symmetric (see Figure 7.24). The performance of both the PS and UPS-tx techniques is essentially unaffected by the number of input images used. From the literature there is a suggestion that illuminant symmetry has a beneficial influence on performance but this was not elaborated upon [Yuille1999]. The sensitivity of the UPS algorithm is likely to stem from the absolute orientation step.



Figure 7.24 Mean TSER for three synthetic textures versus number of images input to three photometric stereo techniques. Illuminant configurations symmetric.
7.3.6 Mean Surface Normal Deviation

Motivation

The UPS-tx algorithm assumes that the surface texture has a planar megastructure. This facilitates the reduction of the ambiguity because the mean surface normal is aligned with the *z*-axis as discussed in Chapter 6. If this was not the case, errors would be likely to be introduced into the orientation stage. It is important to evaluate the performance of the algorithm when the mean surface normal cannot be taken as $[0,0,1]^{T}$.

Data Generation

We created a data set for each of the three synthetic surface textures detailed in Section 3.4 by using Equation 7.1 to generate images of each surface under the range of illumination directions corresponding to: $0^{\circ} \le \tau < 360^{\circ}$, $\Delta \tau = 10^{\circ}$, $30^{\circ} \le \sigma \le 60^{\circ}$, $\Delta \sigma = 15^{\circ}$. Each data set therefore contained 108 images.

Method

Thirty-six images corresponding to a constant illumination slant angle σ of 45° were selected from each data set for use as the input data for the UPS-tx algorithm. Upon each application of the algorithm, an offset was added to the calculated angle between the pseudo mean surface normal and the *z*-axis. We used offset angles of 1°, 2°, 5°, 10° and 15°. The resulting *p* and *q* maps and albedo image were used in conjunction with Equation 7.1 to produce seventy-two relit images with illumination directions corresponding to: $0^{\circ} \le \tau < 360^{\circ}$, $\Delta \tau = 10^{\circ}$, $\sigma = 30^{\circ}/60^{\circ}$. A mean *TSER* value for the three synthetic textures was calculated for each configuration using Equations 3.9-3.11.



Figure 7.25 Mean TSER for three synthetic textures versus mean surface normal offset angle for the UPS-tx algorithm.

Results

It is evident that the performance of the UPS-tx algorithm is adversely affected when the mean surface normal is increasingly offset from the *z*-axis (see Figure 7.25). This is not surprising because this assumption is fundamental to the orientation stage of the algorithm. The implication is that if the UPS-tx algorithm is to produce accurate results, the extent of any offset should be minimised. At the very least this entails ensuring that the texture plane is parallel to the camera sensor during image capture. Setting up the equipment properly is vital in this regard. Unfortunately doing so does not guarantee a mean surface normal of $[0,0,1]^T$, however. This is because the texture mesostructure may also have an influence. The extent to which this is the case depends on the roughness of the surface texture and the chosen sample size. Although this issue has not been addressed by experiment, we may deduce that the mean surface normal of a small section of rough texture is less likely to be aligned with the *z*-axis due to local variation in surface relief. A sufficiently large sample size is therefore required to avoid this problem in order to generate an accurate surface representation of a rough texture using the UPS-tx algorithm.

7.4 Summary & Discussion

In this chapter we evaluated the performance of the UPS-tx algorithm in simulation by using the synthetic textures introduced in Section 3.4. We investigated the effect of shadowing, specular reflection, nearby point illumination and the presence of noise in the input data. We assessed the impact of varying the number of input images and the corresponding illumination configuration. We also evaluated the sensitivity of the algorithm to offsets in the mean surface normal because this is a key assumption in the orientation stage. We tested several other variants of the Lambertian photometric stereo algorithm under identical conditions for comparison.

For Lambertian textures the relight accuracy was found to decrease with increasing surface roughness. This is due to the presence of shadows. A deficiency in the investigation is that only self shadows are considered. Cast shadows were not modelled because accurate rendering was only possible for a limited range of illumination directions. The results presented therefore provide a best case scenario for the presence of shadows in the input data. Nonetheless we may conclude that it is advisable to minimise their presence. This can be done by compensating with favourable lighting positions for extremely rough surfaces by reducing the slant angle. However, the reality is that the method is not well suited to such textures.

For textures which we generated with the Phong model but which were relit with the Lambertian model, relight accuracy was found to decrease for increasing levels of specular reflection. We found that the UPS-tx algorithm is robust for textures with moderate levels of specular reflection. We observed that the breadth of the specular peak is a secondary factor with regard to performance. A narrow peak implies that the resulting highlight is less likely to be observed; a broad peak means that the resulting highlight is effectively diluted. In both cases the resulting reflection behaviour tends towards Lambertian rather than specular and is thus favourable with regard to the accuracy of the UPS-tx technique.

For Lambertian textures we generated with a nearby point source but which were relit assuming a point source at infinity, the accuracy was found to be greater for increased illumination radius. Furthermore, the accuracy was greater for a smaller sample of the texture in question. These findings are reasonable because the effect of nearby point lighting is less pronounced in both cases. In practice, the light source should be positioned as far from the texture as possible. If this is impractical, it may be advisable to implement the algorithm over a series of small sections of a larger surface texture and collate the results. This kind of approach has been used by Rushmeier to determine the bump map of a curved object's surface [Rushmeier1999]. Another approach to this issue is to take a registration image for each illumination direction to adjust the corresponding texture images. However, this is a less practical solution and was not utilised in our experiments.

For Lambertian textures whose images were generated with added noise, accuracy levels were found to steadily decay as the input signal to noise ratio was reduced. An abrupt breakdown in performance was not observed over the range of input noise used. We found that the UPS-tx algorithm is as robust to noise as the other algorithms tested.

The number of input images was found to affect the accuracy of the UPS-tx algorithm but only if the input images were lit in an asymmetric pattern. In the asymmetrical case, more images proved to be beneficial with respect to accuracy. For practical implementation, collecting images of the texture illuminated over a large tilt angle range is to be recommended for favourable performance of this uncalibrated algorithm.

We found that performance is adversely affected when the actual mean surface normal is not aligned with *z*-axis. In order to utilise this algorithm in practice it will therefore be crucial to ensure that the surface texture has a planar megastructure. The sample should be arranged on a level surface and the height variation with regard to the mesostructure of the texture should be relatively small compared to the sample size.

Chapter 8

Assessment of Uncalibrated Photometric Stereo with Texture – Real Data

8.1 Introduction

In the previous chapter we established that the texture-specific algorithm UPS-tx is a viable variant of uncalibrated photometric stereo. In simulation we tested the algorithm under conditions which challenge its assumptions regarding the underlying reflectance model. Images of illuminated synthetic textures were generated to facilitate an assessment of various effects on the algorithm performance. In separate experiments we examined the presence of shadows, specular highlights, nearby point lighting and noise in a controlled manner. With regard to real textures a similar approach would be desirable. However, these effects are influenced rather than precisely controlled by the individual character of each real texture as well as the equipment selection and set-up. Furthermore, it is difficult to consider these deviations in isolation. We decided to centre our analysis of real surface textures on parameters which could be controlled. Our primary investigation therefore examines the effect of illuminant number and configuration on relighting accuracy. We also evaluate the relative performance for the thirty-one real textures and discuss both exceptional and poor performance in terms of the corresponding character of the texture

This chapter is organised as follows:

The techniques utilised in the real texture experiments are described in Section 8.2. In Section 8.3 the performance of the texture-specific technique UPS-tx is assessed. We consider the effect of the number of input images and the corresponding illuminant configuration. The impact of the reflection characteristics of specific textures is then discussed with a view to explaining the observed performance. The findings are summarised and conclusions are drawn in Section 8.4.

8.2 Experimental Technique Variations

In Chapter 7 we utilised five variants of Lambertian photometric stereo (PS-3, PS, UPS, UPS-tx, UPS-it) for the simulation experiments. Despite the extra expense in terms of run time, the iterative technique UPS-it was not found to produce an improvement in relighting accuracy and is not considered further in this thesis. The three-image calibrated technique PS-3 was investigated with real textures in Chapter 4. The results, which are presented in Section 4.6, are not included in this chapter. We decided to focus on a comparison of the algorithms which use the equivalent amount of input data i.e. the PS, UPS and UPS-tx algorithms.

PS

This is the over-constrained version of photometric stereo [Woodham1980] which was described in Section 2.5.3. With regard to relighting accuracy this calibrated technique was consistently found to correspond to the best performance in simulation.

UPS

This is the technique originally proposed by Hayakawa [Hayakawa1994] which is described in Section 6.2. In general UPS was found to correspond to the poorest performance in the simulation experiments.

UPS-tx

This is the texture-specific uncalibrated technique which we developed in Chapter 6. Its performance in simulation was comparable to that for the calibrated technique PS.

8.3 Assessment of Technique Performance

Our goal is to evaluate the UPS-tx algorithm with real texture data and compare its performance to that of the calibrated technique PS and the published uncalibrated technique UPS using relight accuracy as a gauge of performance. With regard to real texture data, the parameters which are readily controlled and hence lend themselves to the design of experiments are illuminant number and configuration. The effect of these factors on relighting accuracy is considered in Section 8.3.1. In Section 8.3.2 we examine performance results in terms of individual textures as opposed to a mean performance over the thirty-one textures. The objective is to facilitate an evaluation of algorithm performance in terms of texture character. In this case we are mainly concerned with shadows and specular highlights whose presence is influenced by the choice of real texture. In Section 8.3.3 we discuss the impact of equipment selection and set-up on algorithm performance since this influences the extent of nearby point lighting effects, noise and mean surface vector deviation. We re-assess results from Sections 8.3.1 & 8.3.2 and compare the effect of different sizes of texture sample with a view to establishing the effect of point lighting.

8.3.1 Illuminant Number & Positioning

Motivation

It is evident from the simulation experiments in Chapter 7 that the relative position of the illumination directions is a potentially significant factor with regard to relight accuracy for the UPS-tx algorithm. We found that the algorithm performance does not actually deteriorate if the corresponding illumination configuration is symmetrical even with a relatively small number of input images. In the case of an asymmetrical configuration, however, we determined that relighting accuracy depends on the number of images input. The UPS-tx algorithm was found to be the most sensitive in this regard compared to the other techniques tested i.e. UPS and PS. It is desirable to repeat these experiments with real surface textures to verify these results.

Data Generation

Colour images of thirty-one real textures were captured using a CCD camera as described in Section 2.5. The data set for each texture consists of 108 images under the range of illumination directions specified in Section 3.4.2. In terms of illumination angles the images correspond to: $0^{\circ} \le \tau < 360^{\circ}$ with $\Delta \tau = 10^{\circ}$ and $30^{\circ} \le \sigma \le 60^{\circ}$ with $\Delta \sigma = 15^{\circ}$. The images which are 1280 × 1024 pixels in size were cut to 128 × 128 pixels to provide input data for the initial experiments in this section.

Method

As in simulation thirty-six images corresponding to a constant illumination slant angle σ of 45° were selected from each data set for use as the *initial* input data for the asymmetric investigation. Subsequent applications of each algorithm (PS, UPS, UPS-tx) involved monotonically reducing the number of input images such that the tilt angle range decreased by 30° each time (0° $\leq \tau \leq 360^{\circ}$ -k30° for k=1-8, $\Delta \tau$ =10°). For the symmetric investigation thirty-six images corresponding to a constant illumination slant angle σ of 45° were again selected from each data set for use as the *initial* input data. Subsequent applications of each algorithm involved monotonically reducing the number of input images such that the illuminant configuration was maintained in a symmetric pattern. In each case the resulting p and q maps and albedo image were used in conjunction with Equation 7.1 to produce seventy-two relit images with illumination directions corresponding to: 0° $\leq \tau < 360^{\circ}$, $\Delta \tau = 10^{\circ}$, $\sigma = 30^{\circ}/60^{\circ}$. A mean texture signal to relight error ratio (*TSER*) value for the thirty-one real textures was calculated for each configuration using Equations 3.9-3.11.



Figure 8.1 Mean TSER versus number of images input to three photometric stereo techniques. Mean TSER averaged over 31 real textures. Illuminant configurations largely asymmetric.

Results

It is apparent that reducing the number of input images such that the corresponding illumination configuration is increasingly asymmetric is detrimental to the performance of the three algorithms for the thirty-one real textures tested (see Figure 8.1). Overall the calibrated over-constrained algorithm PS outperforms the uncalibrated techniques in terms of relight accuracy whilst the UPS-tx algorithm attains consistently better performance than the UPS algorithm. In this case UPS-tx is less sensitive to the number of input images compared to the results reported in simulation in which it was outperformed by the UPS algorithm for lower numbers of input images.



Figure 8.2 Mean TSER versus number of images input to three photometric stereo techniques. Mean TSER averaged over 31 real textures. Illuminant configurations symmetric.

With regard to the symmetric experiments, the performance of the three algorithms is essentially unaffected by a reduction in the number of input images if the corresponding illumination configuration is symmetric (see Figure 8.2). This behaviour was also observed in the simulation experiments. The implication is that even if the illumination directions are unknown, it is advisable to use an even sampling of the illumination hemisphere in order to achieve reasonable levels of accuracy with these algorithms. Unfortunately it is therefore unlikely that the idea of simply "waving a torch randomly to reveal the Euclidean structure" of a surface texture, to paraphrase Drbohlav [Drbohlav2002], will result in sufficient accuracy with the uncalibrated techniques.

8.3.2 Texture Character

Motivation

A notable feature of the results presented in the previous section is that the values observed for our overall measure of accuracy, the mean *TSER*, are significantly lower compared to those in the simulation experiments (compare Figures 7.22, 7.23 with 8.1, 8.2). We attribute this to the fact that Lambert's law was used to relight the synthetic textures for the simulation experiments. Since the three photometric stereo algorithms assume Lambertian behaviour, relatively poor performance is more likely to be observed for real surface textures whose reflectance deviates significantly from the ideal. In simulation we found this to be the case for synthetic surface textures with high rms roughness and for those exhibiting a high proportion of specular reflection. It would be useful to determine if easily discernable texture characteristics such as the presence of shadows and specular highlights correlate with the performance of the Lambertian photometric stereo algorithms for real texture data.

Data Generation

Images of the thirty-one real textures which are 1280×1024 pixels in size were cut to 512×512 pixels to provide data for the experiments in this section. These larger images were generated since the emphasis is on relighting performance with respect to individual textures. As before the data set for each texture consisted of 108 images under the range of illumination directions specified in Section 3.4.2.

Method

We generated two input data sets for each of the thirty-one textures. One contained eighteen symmetrically lit images such that $0^{\circ} \le \tau < 360^{\circ}$ with $\Delta \tau = 20^{\circ}$; the other contained eighteen asymmetrically lit images such that $0^{\circ} \le \tau < 180^{\circ}$ with $\Delta \tau = 10^{\circ}$. A constant slant angle of 45° was used in each case. Each data set was processed with each of the three algorithms (PS, UPS, UPS-tx). The resulting *p* and *q* maps and albedo image were used in conjunction with Equation 7.1 to produce seventy-two relit images with illumination directions corresponding to: $0^{\circ} \le \tau < 360^{\circ}$, $\Delta \tau = 10^{\circ}$, $\sigma = 30^{\circ}/60^{\circ}$. A texture signal to relight error ratio (*TSER*) value was calculated for both configurations for each of the thirty-one real textures using Equations 3.9-3.10.

Results

The results are presented in bar chart form since the individual *TSER* values rather than the mean are of interest in this case (see Figures 8.3 & 8.4). A comparison of the two charts re-emphasises the benefit of using symmetrically-lit input data since relight accuracy is generally greater for this configuration. The significant variation in performance from texture to texture in each chart is of particular note in this case.

(a)



Figure 8.3 TSER (texture signal to relight error ratio) for real textures with input images which are (a) symmetrically lit and (b) asymmetrically lit.

Organising the *TSER* data used to generate Figure 8.3a,b in descending order allows the relative performance ranking of the thirty-one real textures to be evaluated more readily (see Table 8.1). Although the rankings are not identical in each case, it is apparent that *in general* some textures correspond to better performance than others. An example of this is texture g. Although its ranking drops for the UPS technique it corresponds to the best performance for the PS and UPS-tx algorithms with both illumination configurations. The textures *i*, *l*, and *t* appear in the first quartile in each case except for one when they are in the second quartile. At the lower end of the rankings, texture *h* corresponds to one of the worst performances whilst the textures *c*, *v* and *z* mostly feature in the fourth quartile.

Having observed that the choice of texture has a definite bearing on relighting accuracy for the algorithms, it is important to establish a reason for the dependency. We initially attempted to pursue this in an analogous manner to the simulation experiments by, for example, plotting TSER values against the corresponding rms height values for each real texture presented in Chapter 3. This did not provide meaningful Opting for an alternative approach which does not rely on results, however. approximate estimates of relevant parameters, we instead arranged images of the thirty-one real textures according to their ranking (see Figure 8.4); here we focus on the results for the UPS-tx algorithm because it is this technique which has been developed in this thesis. With regard to the images in Figure 8.4, it is difficult to characterise the textures ranked in the second and third quartiles since common features are not apparent. A comparison between the first and fourth quartiles is more enlightening, however. In general the textures of the first quartile appear to have less significant height variation with regard to their mesostructure and are therefore less prone to shadowing than the textures in the fourth quartile. As a result the textures of the fourth quartile are relatively dark in comparison although this is also accounted for in some instances by relatively low albedo (see textures v, c). Textures ac and h are also relatively dark but in addition their images indicate the presence of specular highlights. Specular reflection is not evident in the images of the textures in the first quartile.

Overall, the implication is that deviation from Lambertian behaviour due to the character of the texture is a likely candidate with regard to providing an explanation for the variation in performance. We investigate this further by examining four textures in greater detail. Textures g, ac, ae and h correspond to the maximum, the median, the lowest third quartile ranking and the minimum, respectively. Surface representations

(p, q, albedo) and sample relit images for all thirty-one real textures are provided in Appendix A.

	PS		UPS-tx		UPS	
	symmetric	asymmetric	symmetric	asymmetric	symmetric	asymmetric
Maximum	g	g	g	g	р	р
1 st quartile	aa	aa	aa	i	f	t
	1	t	t	1	g	n
	t	1	i.	е	t	L. L.
	i	е	1	m	aa	i
	е	u	е	р	n	f
	b	m	k	b	i i	ab
	р	р	m	n	0	b
2 nd quartile	m	i	b	t	b	m
	u	ac	q	f	k	е
	k	b	n	q	I	q
	n	k	u	У	r	g
	q	q	f	S	z	k
	ac	n	У	j	m	ο
	f	f	р	x	q	aa
Median	У	0	ac	k	ac	ad
3 rd quartile	ab	S	ο	u	j	j
	S	У	d	w	е	r
	x	d	S	ab	ad	x
	d	ad	j	d	У	w
	ο	ab	а	r	s	У
	w	ae	r	ο	а	S
	j	r	ae	aa	d	u
4 th quartile	v	j	ad	ae	v	z
	r	V	Z	а	ae	d
	а	w	v	ac	С	а
	ae	Z	ab	ad	h	ae
	ad	X	X	V	u	ac
	Z	а	W	Z	ab	V
	С	С	С	h	X	h
Minimum	h	h	h	С	W	С

Table 8.1 Relative performance ranking of the real textures for the photometric stereo algorithms PS, UPS-tx and UPS with either symmetric or asymmetric illumination configuration. Grey background signifies TSER < 10dB.



Figure 8.4 Relative performance ranking of the real textures for the UPS-tx algorithm corresponding to a symmetric illumination configuration.

Maximum Ranking Texture



Figure 8.5 Surface representations of texture g generated by the UPS-tx algorithm.

The texture which was found to correspond to the best performance for the photometric stereo algorithm UPS-tx is texture g. It is a finely woven textile with a checked pattern (see albedo estimation in Figure 8.5) which at the given scale, does not feature dramatic changes in surface relief except for the raised outline of a few embroidered flowers and some wrinkles in the fabric (see p and q maps in Figure 8.5). Neither specular highlights nor shadows were evident in the images captured over the comprehensive range of illumination conditions used in their capture. These factors are conducive to good performance and this finding is therefore not unexpected.



Figure 8.6 Intensity profile for a specific pixel position (number 1000) in a series of 36 16-bit images of texture g corresponding to changing illumination direction of monotonically increasing tilt angle and constant slant angle. Diffuse fit determined using Saito's method [Saito1996].

A typical intensity profile for a single pixel position for texture g is given in Figure 8.6. This demonstrates that the diffuse reflection approximation is appropriate for this texture. We confirmed this observation by fitting a sinusoid using a method detailed by Saito [Saito1996]. This involves using the surface normal $\mathbf{n} = [n_x, n_y, n_z]^T$ and albedo ρ estimates for the pixel position in the calculation of the parameters α , β and γ in the following equation:

$$i(x, y) = \alpha \sin(\tau + \beta) + \gamma$$

$$where \quad \alpha = \rho \sin(\sigma) \sqrt{n_x^2 + n_y^2} , \quad \beta = \frac{n_x}{\sqrt{n_x^2 + n_y^2}} , \quad \gamma = \rho \cos(\sigma) n_z$$
(8.1)

The parameterised Equation 8.1 was used to generate the diffuse fit estimate presented in Figure 8.6.



Median Ranking Texture

Figure 8.7 Intensity images of texture ac with specular highlights (a) not visible and (b) visible.

Specular highlights were observed for texture *ac* during the image capture process at various illumination directions (see Figure 8.5). This texture which was formed with plastic toy bricks was therefore expected to result in a relatively poor performance. This is true of the UPS-tx algorithm when the illumination configuration is asymmetric but its performance corresponds to the median for a symmetric configuration which is considered here (see Table 8.1). A typical profile for a pixel position in an image of this texture helps to explain the reason for this (see Figure 8.8).

It is apparent that this intensity profile is one that can be approximated by the Phong model which was described in Chapter 2 [Phong1975]. Saito's method was applied in order to determine the parameters of the model [Saito1996]. In this case outlying intensity values are initially discarded such that the remaining data largely correspond to diffuse reflection. A sinusoid is fitted to this data by parameterising Equation 8.1 as previously described. The estimated diffuse profile is then subtracted from the observed intensity profile to leave intensity data which relates to the specular peak only. If the logarithm of this specular intensity data is plotted against the logarithm of the dot

product of the view and illumination vectors, a linear relationship will be observed if the Phong model holds (Equation 2.3). This stems from its specular component which can be written as follows:

$$\log i_{specular} = \log \rho_s + n \log(\mathbf{v} \cdot \mathbf{l})$$
(8.2)

Using a least squares method it is straightforward to determine both the gradient and the y-axis intercept which correspond to the Phong exponent n and the log of the specular coefficient ρ_s respectively. The parameters determined in this way for texture ac for this pixel position are $\rho_s = 0.4$ and n = 30. According to the simulation results reported in Chapter 7, these Phong parameter values correspond to reasonable performance for the UPS-tx algorithm with a symmetric illumination configuration (see Figure 7.10) and this has been found in practice. The implication is that although texture ac exhibits specular reflection, the size of its characteristic peak is such that intensity spikes may not always be observed. The apparent deviation from ideal diffuse reflectance for this texture depends on the illumination directions used in the input data and also those used for relighting to determine the accuracy.



Figure 8.8 Intensity profile for a specific pixel position (number 1000) in a series of 36 16-bit images of texture ac under changing illumination direction of monotonically increasing tilt angle and constant slant angle. Diffuse fit and Phong fit determined using Saito's method [Saito1996].

Lowest 3rd Quartile Ranking Texture



Figure 8.9 Original image of texture as $(\tau=90^\circ, \sigma=60^\circ)$, a relit image $(\tau=90^\circ, \sigma=60^\circ)$ generated by using surface representations estimated by UPS-tx with a symmetric illumination configuration & their difference image.

Texture *ae* consists of two house bricks which have been cemented together (see Figure Its lowly ranking for the UPS-tx algorithm can be explained by a closer 8.9). examination of the surface. It is evident that the cement section of the texture is relatively low compared to the surface of the bricks. This difference in height will mean that the texture is prone to cast shadows under illumination directions which correspond to large slant angles. To examine this in more detail we do not consider a typical profile for the texture since this would correspond to the brick section of the images. We opt instead to plot the intensity profile for a pixel position which corresponds to the cement section (see Figure 8.10). In this case it is apparent that the pixel is illuminated in general when the light is positioned such that its beam points along the cement section. As the light is rotated at constant slant angle, however, the relatively higher brick surface serves to occlude the beam and causes cast shadow such that the pixel intensity value deviates from that predicted by the Lambertian model. Not only will this cause the prediction of the surface representation to be less accurate but the fact that the relighting process only takes self shadow into account compounds the problem. The difference image highlights this problem since the maximum error lies along the cement section (see Figure 8.9).



Figure 8.10 Intensity profile for a pixel position (corresponding to the cement section) in a series of 36 16-bit images of texture as under changing illumination direction of monotonically increasing tilt angle and constant slant angle.

Minimum Ranked Texture

Texture h corresponds to the worst performance for the UPS-tx algorithm among the thirty-one real textures. It is a coarse woven fabric incorporating specularly reflecting threads (see Figure 8.11). Highlights feature heavily in many of its images collected under a wide range of illumination conditions and provide a significant deviation from the ideal diffuse reflection assumed by UPS-tx.



Figure 8.11 Original image of texture h (τ =230°, σ =45°), a relit image (τ =230°, σ =45°) generated by using surface representations estimated by UPS-tx with a symmetric illumination configuration & their difference image.

A sample intensity profile for a pixel position is given in Figure 8.12. It is apparent that in this case, the reflectance behaviour is more complex since *two* specular peaks are apparent. Whilst such behaviour was not prevalent for the majority of pixels and solely diffuse and normal specular reflection were also observed, its presence marks a significant deviation from ideal reflection and helps to explain the poor performance. This split off-specular reflection has previously been reported [Pont2003] and is attributed to the structure of man-made materials such as shiny woven fabrics like texture h. The poor performance can be further explained by the fact that the Lambertian model is used to relight the textures. The implication is that highlights visible in an original image will not be present in the corresponding relit image (see Figure 8.11).



Figure 8.12 Intensity profile for a pixel position (corresponding to non-Lambertian reflection) in a series of 36 16-bit images of texture h under changing illumination direction of monotonically increasing tilt angle and constant slant angle. Diffuse fit determined using Saito's method [Saito1996].

8.3.3 Equipment Set-up

Motivation

In the previous section we argued that the variation in performance of the three photometric stereo algorithms among the thirty-one real textures could be largely attributed to the character of the texture. We demonstrated that deviations from ideal reflectance such as the presence of shadowing and specular highlights correspond to relatively low performance in terms of relighting accuracy. In the simulation experiments detailed in Chapter 7 we determined that other sources of error such as noise and the effect of point lighting have an adverse effect on accuracy. With regard to the UPS-tx algorithm this is also true for a deviation in the mean surface vector. Whilst these equipment-based issues could well contribute to the overall deviation from ideal diffuse reflectance, so far in this chapter we have assumed that their influence is small relative to the effect of texture character. In a well set-up laboratory for image capture this should be the case. However, we endeavour to check the validity of this assumption in this section.

Noise Level Assessment

Using the equipment set-up described in Chapter 3, we captured fifty images of a white sheet of paper. The images were analysed to determine the mean and variance. This allowed us to calculate an approximate value for the input signal to noise ratio. It was found to be 25dB. With reference to the results of the noise experiments in Chapter 7, this level of noise is unlikely to have a large impact on the performance of the algorithms using input data captured by our laboratory system.

Point Lighting Assessment

The algorithms we have assessed in this chapter each assume the use of a point light source at infinity. In reality our light source is at a distance of approximately 0.6m from the centre of the texture. From the simulation experiments we determined that for a given illumination radius, the effect of point lighting was reduced by using a smaller sample of texture due to the fact that this serves to limit the range of illumination directions. Since the real texture images in the database were captured at a constant illumination radius, we decided that comparing the performance of different sample sizes would provide a means of investigating whether point lighting has detrimental effect on performance in this case. To do so we examine the values of the mean *TSER* values from Section 8.3.1 with those of Section 8.3.2. These correspond to a square sample of side which measures less than two centimetres (128×128 pixels) and approximately seven centimetres (512×512 pixels) respectively. See Table 8.2.

Image size (pixel)	mean TSER for PS		mean TSER for UPS		mean TSER for UPS-tx	
	symmetric	asymmetric	symmetric	asymmetric	symmetric	asymmetric
128 × 128	12.4 dB	8.8 dB	6.3 dB	6.1 dB	9.5 dB	6.9 dB
512 × 512	11.9 dB	8.8 dB	7.0 dB	5.9 dB	10.7 dB	6.9 dB
Difference	0.5 dB	0.1 dB	0.7 dB	0.1 dB	1.2 dB	0.0 dB

Table 8.2 Mean TSER (texture signal to relight error ratio) averaged over thirty-one real textures for different sizes of texture sample.

Despite a significant difference in physical sample size, the results do not indicate that the use of a smaller sample size is beneficial. We conclude that nearby point lighting does not have an adverse impact on performance in this case.

Mean Surface Normal Deviation

A key assumption in the UPS-tx concerns the mean surface normal. If it is not aligned with the *z*-axis then errors will be introduced into the output data. We measured the angle between the camera line of sight and the texture mounting plate. It was found to be less than 2° . Whilst this small deviation will adversely affect the relighting accuracy obtained for this algorithm we conclude that taking the simulation results into account, it is not significant in comparison to the effect of texture character.

8.4 Summary & Discussion

In this chapter we evaluated the performance of the UPS-tx algorithm with images of the real textures introduced in Chapter 2. We investigated the effect of altering the number of input images and the corresponding illumination configuration. We compared the performance of the UPS-tx uncalibrated algorithm with the equivalent calibrated algorithm PS and a published uncalibrated algorithm UPS. We examined the results of these experiments with a view to explaining the relative performance observed among the textures.

We found that reducing the number of input images such that the illumination configuration becomes increasingly asymmetric is detrimental to the performance of the UPS-tx algorithm with real textures. The PS algorithm was also found to be dependent on the number of input images but outperformed the UPS-tx algorithm. The UPS algorithm was found to have the worst performance of the three. A symmetrical illumination configuration was found to be beneficial to the performance of all three algorithms. Whilst this is impractical to implement for an uncalibrated photometric stereo algorithm, these results suggest that images of the illuminated texture should be captured over a complete range of illumination tilt angles.

The most significant effect on the UPS-tx algorithm performance concerns the character of the texture. We attribute the range of performance observed for the thirty-one real textures to deviation from the ideal reflectance model assumed by the algorithm. We demonstrated that this deviation largely stems from the character of the texture and that equipment-based issues have less influence in comparison. We noted that the effect of such deviations is amplified by our performance measure which compares original texture images against corresponding relit images generated with the Lambertian model. The fact that neither cast shadows nor specular highlights are rendered compounds the problem. We conclude that the method is not suitable for use with textures which exhibit specular reflection and those with significant surface relief.

Chapter 9

Summary and Conclusion

9.1 Summary

In this thesis we investigated the relighting of real world surface textures under arbitrary illumination conditions. This work was motivated by the desire to enable the photorealism of 3D virtual scenes to be enhanced. We proposed to realise this by facilitating the acquisition of the requisite data to model surface texture. Our objective was to develop an inexpensive yet effective method which would be practical to implement with a view to making the technique accessible to a wider range of users.

Chapter 2

Since we are concerned with the appearance of illuminated surfaces, we initially considered the reflection of light in Chapter 2. We discussed the various approaches commonly utilised to model reflectance. Of particular note are Lambert's law for diffuse reflection and the Phong model for specular reflection both of which feature in this work. With regard to relighting, techniques based on reflectance models were extensively reviewed. We also considered image-based techniques. Using the defined criteria to facilitate selection we identified the Lambertian photometric stereo algorithm as the most suitable approach for this work. This model-based technique uses inexpensive equipment, requires few input images and minimal calibration. In this thesis we investigated both the calibrated and uncalibrated versions of the technique. The latter seemed particularly attractive since it only requires intensity images as input; knowledge of the illumination direction corresponding to each input image is not needed. With regard to the algorithm output, the resulting surface-explicit representation is compact. It consists of a bump map which encodes surface gradient data and the corresponding colour albedo image which defines the diffuse reflectance. Significantly this is highly compatible with computer graphics hardware and software.

Chapter 3

In Chapter 3 we defined the term surface texture as referring to globally flat rough patterned surfaces for the purposes of this thesis. The data sets used to assess the performance of the photometric stereo algorithms were then described. They consist of

images of either synthetic or real surface textures illuminated under a set range of directions. The three synthetic surface models used to generate the former are detailed. The equipment utilised to capture images of the thirty-one real textures is also described. We discussed the characteristics of the utilised textures. The texture signal to relight error ratio was defined as a performance measure.

Chapter 4

In Chapter 4 we detailed our investigation into the optimal placement of the illumination vectors in three-image Lambertian photometric stereo. The optimal performance of this calibrated algorithm had only previously been considered in terms of the illumination slant angle [Woodham1980]. Apart from the need to avoid co-planar illumination arrangements, no guidelines with respect to illumination tilt angle are apparent in the literature. We derived an overall figure of merit based on noise variance and used it to investigate this issue. By minimising it we found that the optimal arrangement corresponds to the case when the three illumination vectors are orthogonal to each other. For illumination arrangements of constant slant angle, we found that the optimal difference between successive tilt angles is 120°. This was verified experimentally by assessing the relight error for the thirty-one real textures. With regard to slant angle, we found that for a smooth surface, the slant angle should be maximised for optimal performance. For rough surfaces the optimal slant angle is approximately 55°. This theoretical treatment ignores the effect of shadowing, however. In reality the optimal slant angle will depend on the roughness of the surface. We also noted that the figure of merit which was developed to facilitate this work can be used to determine whether a given illumination configuration is favourable. We found that McGunnigle's scheme with a difference of 90° in tilt angle was sub-optimal but not significantly so [McGunnigle1998].

Chapter 5

We introduced the uncalibrated photometric stereo algorithm in detail in Chapter 5. We reviewed and summarised the various approaches utilised to reduce and resolve the inherent ambiguity. We highlighted the three key assumptions commonly utilised to constrain the solution. We rejected the use of both the integrability and the consistent viewpoint constraints for this work due to the fact that our textures have rough discontinuous surfaces and are diffusely reflecting materials, respectively. We concluded that the constant light intensity assumption which constrains the illumination vectors to be of equal magnitude could potentially be employed, however.

Chapter 6

In Chapter 6 we developed a texture-specific version of the uncalibrated photometric stereo algorithm. We outlined the mathematical framework to implement the constant light source intensity constraint in order to reduce the ambiguity in the solution. The resulting residual ambiguity was shown to be an orthogonal transformation and we proposed a practical procedure to resolve it. The initial step, which exploits the globally flat nature of the surface textures, involves aligning the mean of the surface normal estimates with the *z*-axis. The second step is a rotation about the *z*-axis and requires a single illumination tilt angle to be known. The third step involves an overall scaling and requires a single illumination slant angle to be known. We indicated that the scaling step is not actually required for relighting purposes but that it was utilised in our experiments to enable the performance of the overall algorithm to be evaluated. It is necessary to determine a unique solution in this case to facilitate a valid comparison with other methods.

Chapter 7

We presented the results of the simulation experiments in Chapter 7 with a view to evaluating the performance of the proposed uncalibrated technique. We examined the robustness of the method to deviations from ideal Lambertian reflectance since this is a fundamental assumption of the algorithm. We considered the effect of noise because this will be an inherent feature in the images of real textures captured by camera. We also assessed the effect of input image number and position. The multiple image (>3) calibrated technique was found to be the most accurate of the algorithms tested. The proposed texture-specific uncalibrated method attained a comparable performance in general. One exception to this is in regard to the number of input images. In this case its relight accuracy was considerably more sensitive but only for asymmetric illumination arrangements. We concluded that although this technique is uncalibrated, the illumination direction should be sampled over the full tilt angle range. Relight accuracy was found to decrease for higher values of rms roughness. This is due to the increasing number of shadows in the input images. Hence we concluded that it is advisable to minimise the presence of shadows by using smaller slant angles for very rough surface textures. Relight accuracy was found to decrease as the proportion of specular reflection increases. These experiments allowed us to conclude that the algorithm can be utilised with moderately glossy surface textures. With point lighting where illumination direction varies over the surface we found that both a smaller surface area and a larger illumination radius resulted in greater relight accuracy. We

concluded that in practice we should position the light as far from the texture sample as possible and avoid processing images corresponding to large areas of surface texture. Relight accuracy was found to decrease as the presence of noise increased. These experiments allowed us to conclude that the level of noise observed in our image capture system was acceptable.

Chapter 8

We assessed the results of the texture-specific algorithm with real data in Chapter 8. Its performance was found to compare favourably with its calibrated equivalent although as in simulation, the latter is more accurate in general. Performance was found to improve with larger numbers of input images. A symmetric arrangement was also found to be beneficial although we recognised that this is impractical to implement if the method is truly uncalibrated. We reviewed the relative performance for the thirty-one real textures and discussed extremes of performance in terms of texture character.

9.2 Conclusion

With regard to the calibrated three-image Lambertian photometric stereo technique we determined that the optimal placement of the illumination vectors corresponds to an orthogonal arrangement. If the illumination slant angle is constrained to be constant we proved that the optimal configuration corresponds to a difference between tilt angles of successive illumination vectors of 120°. This theoretical result was verified by experiment. It is noted that these findings had not previously been reported in the literature and are original. Ignoring shadowing, the optimal slant angle was found to be maximum for smooth surfaces and approximately 55° for rough surfaces. The former finding agrees with Woodham's observation which is based on reflectance maps but he does not differentiate between types of surface. Due to the effect of the presence of shadows it is not possible to verify the latter result.

We developed a technique based on uncalibrated photometric stereo which is practical to implement for rough surface textures with a planar megastructure. For relighting purposes, a single illumination tilt angle corresponding to the illumination direction in one of the input images is simply required thus minimising the requisite calibration of equipment. This technique was found to be as robust to adverse reflectance conditions as the equivalent calibrated technique. We found that the accuracy of these Lambertian methods are dependent on how well the Lambertian model approximates the reflectance of the surface texture.

9.3 Contribution Recap

This thesis makes the following contributions with regard to the three-image Lambertian photometric stereo technique:

- The optimal illumination configuration is an orthogonal arrangement.
- If the illuminations vectors are of common slant angle, the optimal illumination configuration corresponds to a tilt angle separation of 120°.
- McGunnigle's scheme [McGunnigle1998] was shown to be sub-optimal but not significantly so.

With regard to the uncalibrated photometric stereo technique:

- A practical implementation is proposed and tested for specific use with rough surface textures.
- Accuracy levels are comparable to the equivalent calibrated technique.

We note that both of these techniques satisfy the criteria specified in Section 1.3 to a large degree. They are suitable for globally flat diffuse surfaces, images of which can be captured and processed with consumer-level equipment. The capture of input data is practical since a small number of images is required in each case. Furthermore, the procedure for the calibrated technique is now explicit with regard to illumination configuration. The technique based on uncalibrated photometric stereo only requires knowledge of one illumination tilt angle and hence reduces the input data burden significantly. Both algorithms avoid the use of iteration and can be regarded as computationally efficient. Finally, we note that the resulting surface representation (the bump map & albedo colour map) is of low dimension and compatible with programmable graphics cards which facilitate real-time per-pixel relighting.

9.4 Future Work

With regard to calibrated photometric stereo, the optimal illumination configuration was determined through the use of a sensitivity analysis. Ideally we would like to develop a mathematical proof for the orthogonal configuration from first principles in order to confirm this intuitive result. We are also keen to extend this work to examine the case when more than three images corresponding to different illumination vectors are used as input. We observed from experiments with both synthetic and real world textures that a symmetric configuration was beneficial in terms of accuracy compared to an asymmetric arrangement (see Sections 7.3.5 & 8.3.1). However, this is merely an observation and requires further investigation to formally identify the optimal conditions. Another opportunity for future work stems from the fact that our sensitivity analysis does not take the distribution of surface normals into account. We briefly considered this point in Section 4.5.3 by speculating that the optimal orthogonal configuration could be oriented to compensate for certain types of surface with a view to minimising the presence of shadows. It would be valuable to develop our approach to incorporate surface normal distribution and shadowing in order to gauge the effect on the optimal configuration.

Whilst we have proposed and tested an algorithm based on uncalibrated photometric stereo, we noted that in reality the algorithm is only pseudo-uncalibrated because a single illumination tilt angle must be provided. We propose to investigate this technique further with a view to identifying a novel means of resolving the ambiguity without resorting to assumptions such as integrability or the presence of specular highlights. Further work is also required to investigate the effect of mean surface normal deviation from the z-axis. We considered this in Section 7.3.6 to a limited extent but it would be interesting to examine the case when the roughness of the surface is such that the mean vector is not $\begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T$ despite a planar megastructure which is parallel to the camera sensor.

Ultimately we seek to develop methods which would enable a simple desktop PC implementation of photometric stereo. In doing so, we aim to facilitate the generation of bump maps and albedo texture maps of real world textures to enhance photorealistic rendering in mixed reality applications.

Appendix A

Texture-specific Uncalibrated Algorithm (UPS-tx) Results

Texture a

Surface Representations.

1. Symmetric illumination configuration





2. Asymmetric illumination configuration





albedo image



albedo image



Relight comparison. 1. $\tau=20^{\circ} \sigma=30^{\circ}$

Original



2. τ=120° σ=60°

Original



Symmetric





Asymmetric



Asymmetric



$\frac{\text{Texture } b}{\text{Surface Representations.}}$

1. Symmetric illumination configuration



2. Asymmetric illumination configuration



Relight comparison. 1. $\tau=20^{\circ} \sigma=30^{\circ}$





2. τ =120° σ =60°

Original





Symmetric

Symmetric

albedo image



albedo image



Asymmetric



Asymmetric



Texture c **Surface Representations**. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. τ=20° σ=30°





2. τ=120° σ=60°

Original





Symmetric

Symmetric

albedo image



albedo image



Asymmetric



Asymmetric





<u>Texture e</u> Surface Representations.

1. Symmetric illumination configuration





2. Asymmetric illumination configuration





albedo image



albedo image



Relight comparison. 1. $\tau=20^{\circ} \sigma=30^{\circ}$

Original



2. τ=120° σ=60°

Original







Asymmetric



Symmetric



Asymmetric



Texture f Surface Representations. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. τ=20° σ=30°

Original



2.
$$\tau$$
=120° σ =60°

Original





albedo image



albedo image



Asymmetric



Symmetric

Symmetric



Asymmetric



$\frac{\text{Texture } g}{\text{Surface Representations.}}$

1. Symmetric illumination configuration





albedo image

2. Asymmetric illumination configuration











$\frac{\text{Texture } h}{\text{Surface Representations.}}$

1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. $\tau=20^{\circ} \sigma=30^{\circ}$

 $1. \tau = 20^{-} \sigma = 3$ Original



Original





Symmetric

Symmetric

albedo image



albedo image



Asymmetric



Asymmetric



Texture *i* **Surface Representations**. 1. Symmetric illumination configuration









Relight comparison. 1. τ=20° σ=30°

Original



2. τ=120° σ=60°

Original







Symmetric



Asymmetric


<u>Texture</u> j Surface Representations.

1. Symmetric illumination configuration



2. Asymmetric illumination configuration









albedo image



Relight comparison. 1. $\tau = 20^{\circ} \sigma = 30^{\circ}$

Original



2. τ=120° σ=60°

Original





Asymmetric







Texture k **Surface Representations**. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. $\tau = 20^{\circ} \sigma = 30^{\circ}$



2. τ=120° σ=60°

Original







albedo image



Symmetric

Symmetric

1112





Texture l Surface Representations. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration





albedo image



albedo image



Relight comparison. 1. τ=20° σ=30°



2. τ=120° σ=60°

Original





Symmetric



Asymmetric





Texture m Surface Representations. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. $\tau = 20^{\circ} \sigma = 30^{\circ}$



2. τ=120° σ=60°

Original





albedo image



albedo image



Asymmetric



Symmetric

Symmetric





<u>Texture *n*</u> Surface Representations.

1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. $\tau=20^{\circ} \sigma=30^{\circ}$



2. $\tau=120^{\circ} \sigma=60^{\circ}$

Original





albedo image



albedo image



Asymmetric



Symmetric

Symmetric





Texture o Surface Representations.

1. Symmetric illumination configuration





2. Asymmetric illumination configuration





albedo image



Relight comparison. 1. τ=20° σ=30°



2. τ=120° σ=60°









Asymmetric



Asymmetric



Symmetric

Symmetric

Texture p Surface Representations. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. $\tau=20^{\circ} \sigma=30^{\circ}$



2. τ=120° σ=60°







Symmetric



Symmetric



albedo image



albedo image



Asymmetric





$\frac{\text{Texture } q}{\text{Surface Representations.}}$

1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. $1. \tau=20^{\circ} \sigma=30^{\circ}$ Original



2. τ=120° σ=60°

Original





Symmetric

Symmetric

albedo image



albedo image



Asymmetric







$\frac{\text{Texture } r}{\text{Surface Representations.}}$

1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. $1. \tau=20^{\circ} \sigma=30^{\circ}$



2. τ=120° σ=60°

Original





Symmetric

Symmetric

albedo image



albedo image



Asymmetric





<u>Texture</u> <u>s</u> Surface Representations.

1. Symmetric illumination configuration





2. Asymmetric illumination configuration







Relight comparison. $1. \tau=20^{\circ} \sigma=30^{\circ}$ Original



2. τ=120° σ=60°

Original





Symmetric

Symmetric





Asymmetric



Asymmetric



Texture t **Surface Representations**. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. τ=20° σ=30°





Original





albedo image



albedo image



Asymmetric



Symmetric

Symmetric





Texture *u* **Surface Representations**. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. τ=20° σ=30° Original



2. τ=120° σ=60°

Original





Symmetric

Symmetric

albedo image



albedo image



Asymmetric





Texture v **Surface Representations**. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. τ=20° σ=30°



2. τ=120° σ=60°

Original





Symmetric

Symmetric

albedo image



albedo image



Asymmetric







<u>Texture</u> <u>x</u> **Surface Representations**. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. $\tau=20^{\circ} \sigma=30^{\circ}$



2. τ=120° σ=60°

Original





Symmetric



Symmetric



albedo image



albedo image



Asymmetric





Texture y Surface Representations. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. τ=20° σ=30°



2. τ=120° σ=60°

Original





Symmetric

Symmetric

albedo image



albedo image



Asymmetric





$\frac{\text{Texture } z}{\text{Surface Representations.}}$

1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. $1. \tau=20^{\circ} \sigma=30^{\circ}$ Original



Symmetric





2. τ=120° σ=60°

Original







albedo image



albedo image



<u>Texture aa</u> Surface Representations.

1. Symmetric illumination configuration





2. Asymmetric illumination configuration



albedo image

albedo image



Relight comparison.



2. τ=120° σ=60°









Asymmetric





Texture ab Surface Representations. 1. Symmetric illumination configuration





2. Asymmetric illumination configuration *q* map



Relight comparison. 1. τ=20° σ=30° Original





Symmetric





albedo image



albedo image



Asymmetric



Asymmetric



<u>Texture *ac*</u> Surface Representations.

1. Symmetric illumination configuration





2. Asymmetric illumination configuration



q map

albedo image



albedo image



Relight comparison. 1. $\tau=20^{\circ} \sigma=30^{\circ}$





2. τ=120° σ=60°

Original





Symmetric



Asymmetric





<u>Texture</u> <u>ad</u> **Surface Representations**.

1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. $1. \tau=20^{\circ} \sigma=30^{\circ}$



2. τ=120° σ=60°

Original





Symmetric

Symmetric

albedo image



albedo image



Asymmetric





<u>Texture</u> <u>ae</u> Surface Representations.

1. Symmetric illumination configuration





2. Asymmetric illumination configuration



Relight comparison. 1. $\tau=20^{\circ} \sigma=30^{\circ}$



2. τ=120° σ=60°

Original





Symmetric

Symmetric

albedo image



albedo image



Asymmetric





Appendix B

Sensitivity Expressions for Common Slant Angle

The derivation of the sensitivity expressions presented in Chapter 4 simplifies if the slant angle is common to each of the illumination vectors. In this case the illumination matrix equivalent to Equation 4.1 is given by:

	$\cos \tau_1 \sin \sigma$	$\sin \tau_1 \sin \sigma$	$\cos\sigma$
L =	$\cos \tau_2 \sin \sigma$	$\sin \tau_2 \sin \sigma$	$\cos\sigma$
	$\cos \tau_3 \sin \sigma$	$\sin au_3 \sin \sigma$	$\cos\sigma$

Taking the inverse of **L** provides expressions for the scaled surface normal equivalent to Equations 4.2, 4.3 & 4.4:

$$\begin{split} s_x &= -\left(\frac{(\sin\tau_3 - \sin\tau_2)i_1 + (\sin\tau_1 - \sin\tau_3)i_2 + (\sin\tau_2 - \sin\tau_1)i_3}{K\sin\sigma}\right)\\ s_y &= -\left(\frac{(\cos\tau_2 - \cos\tau_3)i_1 + (\cos\tau_3 - \cos\tau_1)i_2 + (\cos\tau_1 - \cos\tau_2)i_3}{K\sin\sigma}\right)\\ s_z &= \left(\frac{\sin(\tau_3 - \tau_2)i_1 + \sin(\tau_1 - \tau_3)i_2 + \sin(\tau_2 - \tau_1)i_3}{K\cos\sigma}\right)\\ \end{split}$$

where $K = \sin(\tau_3 - \tau_2) + \sin(\tau_1 - \tau_3) + \sin(\tau_2 - \tau_1)$

The nine sensitivity expressions equivalent to 4.5 - 4.13 are:

$$\frac{\partial s_x}{\partial i_1} = -\left(\frac{\sin\tau_3 - \sin\tau_2}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_y}{\partial i_1} = -\left(\frac{\cos\tau_2 - \cos\tau_3}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_z}{\partial i_1} = \left(\frac{\sin(\tau_3 - \tau_2)}{K\cos\sigma}\right) \\ \frac{\partial s_x}{\partial i_2} = -\left(\frac{\sin\tau_1 - \sin\tau_3}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_y}{\partial i_2} = -\left(\frac{\cos\tau_3 - \cos\tau_1}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_z}{\partial i_2} = \left(\frac{\sin(\tau_1 - \tau_3)}{K\cos\sigma}\right) \\ \frac{\partial s_x}{\partial i_3} = -\left(\frac{\sin\tau_2 - \sin\tau_1}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_y}{\partial i_3} = -\left(\frac{\cos\tau_1 - \cos\tau_2}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_z}{\partial i_3} = \left(\frac{\sin(\tau_2 - \tau_1)}{K\cos\sigma}\right) \\ \frac{\partial s_z}{\partial i_3} = -\left(\frac{\sin\tau_2 - \sin\tau_1}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_z}{\partial i_3} = -\left(\frac{\cos\tau_1 - \cos\tau_2}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_z}{\partial i_3} = \left(\frac{\sin(\tau_2 - \tau_1)}{K\cos\sigma}\right) \\ \frac{\partial s_z}{\partial i_3} = -\left(\frac{\sin\tau_2 - \sin\tau_1}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_z}{\partial i_3} = -\left(\frac{\sin\tau_2 - \sin\tau_1}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_z}{\partial i_3} = \left(\frac{\sin(\tau_2 - \tau_1)}{K\cos\sigma}\right) \\ \frac{\partial s_z}{\partial i_3} = -\left(\frac{\sin\tau_2 - \sin\tau_1}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_z}{\partial i_3} = -\left(\frac{\sin\tau_2 - \sin\tau_1}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_z}{\partial i_3} = \left(\frac{\sin\tau_2 - \tau_1}{K\cos\sigma}\right) \\ \frac{\partial s_z}{\partial i_3} = -\left(\frac{\sin\tau_2 - \sin\tau_1}{K\sin\sigma}\right) \qquad \qquad \frac{\partial s_z}{\partial i_3} = \left(\frac{\sin\tau_2 - \tau_1}{K\cos\sigma}\right) \\ \frac{\partial s_z}{\partial i_3} = \left(\frac{\sin\tau_2 - \tau_1}{K\cos\sigma}\right) \qquad \qquad \frac{\partial s_z}{\partial i_3} = \left(\frac{\sin\tau_2 - \tau_1}{K\cos\sigma}\right) \\ \frac$$

We note that the slant angle σ features only in the denominator of scaled surface normal and intensity expressions. Since the x and y components take sin σ whilst z takes cos σ . This facilitates an understanding of the different behaviour observed between rough and smooth surfaces. With regard to the latter when the z element is omitted, increasing the slant angle increases the denominator for the six sensitivities hence reducing the overall noise. For a rough surface when we include the z component increasing the slant angle decreases its denominator and offsets the decrease in overall noise for the x and y components.

Appendix C

Statistics Proofs

The following proofs may be found in many statistics textbooks e.g. [Davis1986] but are given here for completeness.

To determine an expression for noise in the scaled surface normal vector s we considered the variance of a parameter x which is the function of two variables u & v.

$$x_i = f(u_i, v_i)$$

Employing a Taylor series expansion, the following expressions are derived:

$$f(u_i, v_i) = f(\overline{u}, \overline{v}) + (u_i - \overline{u}) f_u'(\overline{u}, \overline{v}) + (v_i - \overline{v}) f_v'(\overline{u}, \overline{v}) + \dots$$

If the function is first order with regard to both variables then the mean value of x is: $\overline{x} = f(\overline{u}, \overline{v})$

and the overall expression simplifies:

....

$$x_{i} = \overline{x} + (u_{i} - \overline{u})\frac{\partial x}{\partial u} + (v_{i} - \overline{v})\frac{\partial x}{\partial v}$$
$$x_{i} - \overline{x} = (u_{i} - \overline{u})\frac{\partial x}{\partial u} + (v_{i} - \overline{v})\frac{\partial x}{\partial v}$$

The variance of the parameter *x* is written as follows:

$$\psi_x^2 = \frac{1}{N} \sum_{i=1}^{N} (x_i - \overline{x})^2$$

$$\psi_x^2 = \frac{1}{N} \sum_{i=1}^{N} [(u_i - \overline{u}) \frac{\partial x}{\partial u} + (v_i - \overline{v}) \frac{\partial x}{\partial v}]^2$$

$$\psi_x^2 = \frac{1}{N} \sum_{i=1}^{N} [(u_i - \overline{u})^2 \left(\frac{\partial x}{\partial u}\right)^2 + (v_i - \overline{v})^2 \left(\frac{\partial x}{\partial v}\right)^2 + 2(u_i - \overline{u})(v_i - \overline{v}) \frac{\partial x}{\partial u} \frac{\partial x}{\partial v}]$$

$$\psi_x^2 = \frac{1}{N} \sum_{i=1}^{N} (u_i - \overline{u})^2 \left(\frac{\partial x}{\partial u}\right)^2 + \frac{1}{N} \sum_{i=1}^{N} (v_i - \overline{v})^2 \left(\frac{\partial x}{\partial v}\right)^2 + \frac{1}{N} \sum_{i=1}^{N} 2(u_i - \overline{u})(v_i - \overline{v}) \frac{\partial x}{\partial u} \frac{\partial x}{\partial v}$$

The general equation (Equation 4.14) for the variance of x which is a function of two variables u & v is :

$$\psi_x^2 = \psi_u^2 \left(\frac{\partial x}{\partial u}\right)^2 + \psi_v^2 \left(\frac{\partial x}{\partial v}\right)^2 + 2\psi_{uv}^2 \frac{\partial x}{\partial u} \frac{\partial x}{\partial v}$$

Independent Parameters

If the input parameters are independent the general expression simplifies to Equation 4.15:

$$\psi_x^2 = \psi_u^2 \left(\frac{\partial x}{\partial u}\right)^2 + \psi_v^2 \left(\frac{\partial x}{\partial v}\right)^2$$

Highly Correlated Parameters

To derive the overall figure of merits we assumed that the noise processes corresponding to each component of the scaled surface normal are highly correlated since each are a function of the three image noise processes. To illustrate this case we use the following example:

x = u + v where v = au and *a* is a constant scalar

The variance equations for each of the variables *u* & *v* are written as follows:

$$\psi_u^2 = \frac{1}{N} \sum_{i=1}^{N} (u_i - \overline{u})^2 \qquad \qquad \psi_v^2 = \frac{1}{N} \sum_{i=1}^{N} (v_i - \overline{v})^2 \qquad \qquad \psi_v^2 = \frac{1}{N} \sum_{i=1}^{N} (au_i - a\overline{u})^2 \qquad \qquad \qquad \psi_v^2 = a^2 \frac{1}{N} \sum_{i=1}^{N} (u_i - \overline{u})^2$$

It is apparent that the following expressions are true:

1.
$$\psi_v^2 = a^2 \psi_u^2$$

2. $\frac{\partial x}{\partial u} = \frac{\partial x}{\partial v} = 1$

Using the general formula, the following can then be deduced:

$$\psi_x^2 = \psi_u^2 + a^2 \psi_u^2 + 2a \psi_u^2$$
$$\psi_x^2 = \psi_u^2 (a+1)^2$$
$$\psi_x = \psi_u + a \psi_u$$
$$\psi_x = \psi_u + \psi_v$$

This result allowed us to derive the figures of merit given by Equations 4.23 & 4.27.

Appendix D

Optimal Performance

In Chapter 4 we presented four plots which were relevant to the discussion. Here we present further plots for the sake of completeness.

Tilt Angle τ_3 Variation

Here we present the plots for the smooth surface figure of merit M_{smooth} in addition to those for M_{rough} (Figures 4.4 & 4.5) for comparison.



• Figure of merit versus third tilt angle τ_3 with $\tau_1=0^\circ$, $\tau_2=120^\circ$.

• Figure of merit versus third tilt angle τ_3 with $\tau_1=0^\circ$, $\tau_2=90^\circ$.



We draw the same conclusions from the M_{smooth} plots as from those of M_{rough} . A minimum value for the figure of merit is apparent in each case. This corresponds to $\tau_3=240^\circ$ when $\tau_1=0^\circ$, $\tau_2=120^\circ$. When $\tau_1=0^\circ$, $\tau_2=90^\circ$ it is apparent that McGunnigle's simplified photometric scheme [McGunnigle1998] is sub-optimal but not significantly so.

Slant Angle σ Variation

Here we present the plots for McGunnigle's scheme ($\Delta \tau$ =90°, common slant angle) for comparison with the those corresponding to the optimal configuration ($\Delta \tau$ =120°, common slant angle) which are Figures 4.7 & 4.8.

• Rough Surface

(i) Δ*τ*=90°

(ii) $\Delta \tau = 120^{\circ}$



Again, we draw the same conclusions with the results for $\Delta \tau = 90^{\circ}$. A minimum figure of merit value is apparent for the rough surface and corresponds to a slant angle of ~55°. When the surface is smooth the slant angle corresponding to the minimum tends towards 90°.

Appendix E

List of Publications by the Author

Optimal illumination for three-image photometric stereo acquisition of surface texture.

A. D. Spence and M.J. Chantler

Abstract

The optimal placement of the illumination for three-image photometric stereo acquisition of smooth and rough surface textures is derived and verified experimentally. The sensitivities of the scaled surface normal elements are derived and used to provide expressions for the noise variances. An overall figure of merit is developed by considering image-based rendering (i.e. relighting) of Lambertian surfaces. This metric is optimised with respect to the illumination angles. The optimal separation between the tilt angles of successive illumination vectors was found to be 120°. The optimal slant angle was found to be 90° for smooth surface textures and 55° for rough surface textures.

Refereed paper published in the Proceedings of the 3rd International Workshop on Texture Analysis and Synthesis, pp. 89-94, October 2003.

On capturing 3D isotropic surface texture using uncalibrated photometric stereo.

A. D. Spence and M.J. Chantler

Abstract

We propose an uncalibrated method for acquiring the normal and albedo fields of an isotropic 3D surface texture illuminated at a constant slant angle. The method is 'uncalibrated' in that the illumination vectors are not known a priori. We assume single point lighting of a rough Lambertian surface lying in the x-y plane. We use Hayakawa's uncalibrated photometric stereo algorithm to simultaneously estimate the scaled surface normals and the illumination vectors in an arbitrary co-ordinate system. The use of constant illumination slant intensity data means that the required orientation to a viewer co-ordinate system simply involves a z-axis rotation. Orientation in the x-y plane is determined by applying a frequency domain method for estimating illumination tilt angles. Preliminary results from simulations and real data are provided.

Refereed paper published in the Proceedings of the 3rd International Workshop on Texture Analysis and Synthesis, pp. 83-87, October 2003.

Real-time per-pixel rendering of textiles for virtual textile catalogues.

A. Spence, M. Robb, M. Timmins and M.J. Chantler.

Abstract

We present recent results from an EPSRC funded project VirTex (Virtual Textile Catalogues). The goal of this project is to develop graphics and image-processing software for the capture, storage, search, retrieval and visualisation of 3D textile samples. The ultimate objective is to develop a web-based application that allows the user to search a database for suitable textiles and to visualize selected samples using real-time photorealistic 3D animation.

The main novelty of this work is in the combined use of photometric stereo and realtime per-pixel-rendering for the capture and visualisation of textile samples. Photometric stereo is a simple method that allows both the bump map and the colour map of a surface texture to be captured digitally. It uses a single fixed camera to obtain three images under three different illumination conditions. The colour map is the image that would be obtained under diffuse lighting. The bump map describes the small undulations of the surface relief. When imported into a standard graphics program these images can be used to texture 3D models. The appearance is particularly photorealistic, especially under changing illumination and viewpoints. The viewer can manipulate both viewpoint and lighting to gain a deeper perception of the properties of the textile sample. In addition, these images can be used with 3D models of products to provide extremely accurate visualisations for the customer.

Until recently these images could only be rendered using ray-tracing software. However, recent consumer-level graphics cards from companies such as Nvidia, ATI and 3Dlabs now provide real-time per-pixel shading. We have developed software that takes advantage of the advanced rendering features of these cards to render images in real-time. It uses photometrically acquired bump and colour maps of textiles to provide real-time visualisation of a textile sample, under user-controlled illumination, pose and flex.

Published in the International Journal of Clothing Science and Technology, 16(1), 51-62, 2004.

References

Barsky2001	S. Barsky and M. Petrou. Colour photometric stereo : simultaneous reconstruction of local gradient and colour of rough textured surfaces. In <i>Proceedings of International Conference on Computer Vision</i> , pp. 600-605, 2001.
Basri2001	R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. In <i>Proceedings of the</i> δ^{th} <i>IEEE International Conference on Computer Vision</i> , 2, pp.383-390, 2001.
Bayer1976	B.E. Bayer. Color imaging array. U.S. Patent No. 3,971,065, 1976.
Belhumeur1998	P. Belhumeur and D. Kriegman. What is the set of images of an object under all possible lighting conditions? <i>International Journal of Computer Vision</i> , 28(3), pp.245-260, 1998.
Belhumeur1999	P. N. Belhumeur, D. J. Kriegman and A. L. Yuille. The bas- relief ambiguity. <i>International Journal of Computer Vision</i> , 35(1), 1999.
Blinn1977	J.F. Blinn. Models of light reflection for computer synthesized pictures. In <i>Proceedings of Computer Graphics</i> , SIGGRAPH, pp. 192-198, 1977.
Blinn1978	J. Blinn. Simulation of wrinkled surfaces. In Proceedings of Computer Graphics, SIGGRAPH 78, pp. 286-292, August 1978.
Cabral1987	B. Cabral, N. Max and R. Springmeyer. Bidirectional reflection functions from surface bump maps. <i>Computer Graphics</i> , 21(4), pp. 273-281, July 1987.
Chantler1995	M.J. Chantler. Why illuminant direction is fundamental to texture analysis. <i>IEE Proceedings of Vision, Image and Signal Processing</i> , 142(4), pp.199-206, August 1995.
Chantler1997	M. J. Chantler and G. B. Delguste. Illuminant-tilt estimation from images of isotropic texture. <i>IEE Proceedings, Vis. Im. Sig. Proc.</i> , 144(4), 1997.
Cho1993	C. Cho and H. Minamitani. A new photometric method using 3 point light sources. <i>IEICE Transactions on Informations and Systems</i> , E76-D(8), pp. 898-904, 1993.
Chojnacki1994	W. Chojnacki and M.J. Brooks. Revisiting Pentland's estimator of light source direction. <i>Journal of the Optical Society of America A</i> , 11(1), January 1994.
Christensen1994	P.H. Christensen and L.G. Shapiro. Three-dimensional shape from colour photometric stereo. <i>International Journal of Computer Vision</i> , 13(2), pp. 213-227, 1994.
Clark1992	J.J. Clark. Active photometric stereo. In <i>Proceedings of IEEE Conference on Computer Vision and Pattern Recognition</i> , 1, pp. 29-34, 1992.

Clark1999	J.J. Clark and H. Pekau. An integral formulation for differential photometric stereo. In <i>Proceedings of IEEE Conference on Computer Vision and Pattern Recognition</i> , 1, pp. 124-134, 1999.
Coleman1982	E.N. Coleman and R. Jain. Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. <i>Computer Graphics and Image Processing</i> , 18, pp. 309-328, 1982.
Cook1982	R.L.Cook and K.E. Torrance. A reflectance model for computer graphics. <i>ACM Transactions on Graphics</i> , 1(1), 7-24, January 1982.
Dana1999	K.J. Dana, B. van Ginneken, S.K. Nayar and J.J. Koenderink. Reflectance and texture of real-world surfaces. <i>ACM Transactions on Graphics</i> , 18(1), 1999.
Davis1986	J.C. Davis. Statistics and Data Analysis in Geology. 2 nd Edition. Wiley, 1986.
Debevec2000	P. Debevec, T. Hawkins, C. Tchou, H. Duiker, W. Sarokin and M. Sagar. Acquiring the reflectance field of a human face. In <i>Proceedings of Computer Graphics, SIGGRAPH 2000</i> , pp 145-156, 2000.
Dong2002	J. Dong and M.J. Chantler. Capture and synthesis of 3D surface texture. In <i>Proceedings of the</i> 2^{nd} <i>International Workshop on Texture Analysis and Synthesis</i> , pp. 41-45, June 2002.
Dong2003a	J. Dong and M.J. Chantler. Comparison of five 3D surface texture synthesis methods. In <i>Proceedings of the</i> 3^{rd} <i>International Workshop on Texture Analysis and Synthesis</i> , pp. 19-24, October 2003.
Dong2003b	J. Dong. Three-dimensional surface texture synthesis. <i>Ph.D. thesis</i> , School of Mathematics and Computer Sciences, Heriot-Watt University, September 2003.
Drbohlav2001	O. Drbohlav and R. Šára. Unambiguous determination of shape from photometric stereo with unknown light sources. <i>International Conference of Computer Vision</i> , 2001.
Drbohlav2002	O. Drbohlav and R. Šára. Specularities reduce ambiguity of uncalibrated photometric stereo. <i>European Conference of Computer Vision</i> , pp. 46-60, 2002
Drbohlav2003	O.Drbohlav. Towards uncalibrated photometric stereo for non- Lambertian surfaces. <i>Ph.D. thesis</i> , Center for Machine Perception, Department of Cybernetics, Faculty of Electrical Engineering, Czech Technical University, Prague, 2003.
Epstein1995	R. Epstein, P.W. Hallinan and A.L. Yuille. 5±2 Eigenimages suffice: an empirical investigation of low-dimensional lighting models. In <i>Proceedings of the Workshop on Physics-Based Modeling in Computer Vision</i> , pp. 108-116, June 1995.

Fan1997 J. Fan and L.B. Wolff. Surface curvature and shape reconstruction from unknown multiple illumination and integrability. Computer Vision and Image Understanding, 65(2), pp. 347-359, February 1997. Frankot1988 R.T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. IEEE Transactions on Pattern Analysis and Machine Intelligence, 10(4), July 1988. Georghiades1999 A.S. Georghiades, P.N. Belhumeur and D.J. Kriegman. Illumination-based image synthesis: creating novel images of human faces under differing pose and lighting. In Proceedings of the Workshop on Multi-view Modeling and Analysis of Visual Scenes, pp. 47-54, 1999. Georghiades2003a A.S. Georghiades. From few to many : generative appearancebased models for face recognition. *Ph.D. thesis*. Department of Electrical Engineering, Yale University, New Haven, U.S.A, 2003. Georghiades2003b A.S. Georghiades. Recovering 3-D shape and reflectance from a small number of photographs. Eurographics Symposium on Rendering, 2003. Georghiades2003c A.S. Georghiades. Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In Proceedings of the 9th International Conference on Computer Vision, 2003. Gullón2002 C. Gullón. Height recovery of rough surfaces from intensity images. Ph.D. thesis. School of Engineering and Physical Sciences, Heriot-Watt University, 2002. Hallinan1994 P.W. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. In Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, pp. 995-999, 1994. Han2003 J.Y. Han and K. Perlin. Measuring birdirectional texture reflectance with a kaleidoscope. ACM Transactions on Graphics (TOG), 22(3), pp.741-748, July 2003. Hanrahan1993 P. Hanrahan and W. Krueger. Reflection from layered surface due to subsurface scattering. In Proceedings of 20th Annual Conference on Computer Graphics and Interactive Techniques, pp. 165-174, 1993. P. Hansson and P-Å. Johansson. Topography and reflectance Hansson2000 analysis of paper surfaces using a photometric stereo method. Optical Engineering, 39(9), pp. 2555-2561, September 2000. Hayakawa1994 H. Hayakawa. Photometric stereo under a light source with arbitrary motion. Journal of the Optical Society of America A, 11(11), 1994.

He1991	X.D. He, K.E. Torrance, F.X. Sillion and D.P. Greenberg. A comprehensive physical model for light reflection. <i>Computer Graphics</i> , 25(4), July 1991.
Hertzmann2003	A. Hertzmann and S.M. Seitz. Shape and materials by example: a photometric stereo approach. In <i>Proceedings of Computer</i> <i>Vision and Pattern Recognition</i> , 1, pp. 533-540, 2003.
Horn1986	B.K.P. Horn. Robot Vision. MIT Press, Cambridge, MA, 1986.
Horn1987	B. K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. <i>Journal of the Optical Society of America</i> A , 4(4), 1987.
Horn1989	B.K.P. Horn and M.J. Brooks. Shape from shading. MIT Press, Cambridge, MA, 1989.
Ikeuchi1981	K. Ikeuchi. Determining surface orientation of specular surfaces by using the photometric stereo method. <i>IEEE Transactions on</i> <i>Pattern Analysis and Machine Intelligence</i> , 3(6), pp. 661-669, November 1981.
Ikeuchi1991	K. Ikeuchi and K. Sato. Determining reflectance parameters using range and brightness images. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 13(11), pp. 1139-1153, November 1991.
Iwahori1990	Y. Iwahori, H. Sugie and N. Ishii. Reconstructing shape from shading images under point illumination. In <i>Proceedings of 10th International Conference on Pattern Recognition</i> , pp. 83-87, 1990.
Iwahori1992	Y. Iwahori, R.J. Woodham and N. Ishii. Shape from shading with a nearby moving light source. In <i>Proceedings of 2^{nd} International Conference on Automation</i> , Robotics and Computer Vision, 5(5), pp.1-5, 1992.
Iwahori1994	Y. Iwahori, R.J. Woodham, H. Tanaka and N. Ishii. Moving point light source photometric stereo. <i>IEICE Transactions on Information and Systems</i> , E77-D(8), pp. 925-929, 1994.
Jiang1991	X.Y. Jiang and H. Bunke. On error analysis for surface normals determined by photometric stereo. <i>Signal Processing</i> , 23, pp. 221-226, 1991.
Kay1995	G. Kay and T. Caelli. Estimating the parameters of an illumination model using photometric stereo. <i>Graphical Models and Image Processing</i> , 57(5), pp. 365-388. September 1995.
Kim1997	B-H. Kim and R-H. Park. Shape from shading and photometric stereo using surface approximation by Legendre polynomials. <i>Computer Vision and Image Understanding</i> , 66(3), pp. 255-270, 1997.
Klette1999	R. Klette, R. Kozera and K. Schlüns. Reflectance-based shape recovery. In <i>Handbook of Computer Vision and Applications</i> , Vol. 2, Academic Press 1999, pp. 531-590, 1999.

Knill1990	D. C. Knill. Estimating illuminant direction and degree of surface relief. <i>Journal of the Optical Society of America A</i> , 7(4), pp 759-775, 1990.
Koenderink1996a	J.J. Koenderink, A.J.van Doorn and M. Stavridi. Bidirectional reflection distribution function expressed in terms of surface scattering modes. In <i>Proceedings of European Conference on Computer Vision</i> , 2, pp.28-39, 1996.
Koenderink1996b	J.J. Koenderink and A.J. van Doorn. Illuminance texture due to surface mesostructure. <i>Journal of the Optical Society of America</i> A, 13(3), pp. 452-463, March 1996.
Koenderink1997	J. J. Koenderink and A. J. Van Doorn. The generic bilinear calibration-estimation problem. <i>International Journal of Computer Vision</i> , 23(3), pp. 217-234,1997.
Koenderink2003	J.J. Koenderink and S.C. Pont. Irradiation direction from texture. <i>Journal of the Optical Society of America</i> , 20(10), 1875-1882, 2003.
Koudelka2001	M.L. Koudelka, P.N. Belhumeur, S. Magda and D.J. Kriegman. Image-based modeling and rendering of surfaces with arbitrary BRDFs. In <i>Proceedings of IEEE Conference on Computer</i> <i>Vision and Pattern Recognition</i> , 1, pp. 8-14, 2001.
Koudelka2003	M.L.Koudelka, S. Magda, P.N. Belhumeur and D.J. Kriegman. Acquisition, compression and synthesis of bidirectional texture functions. In <i>Proceedings of the</i> 3^{rd} <i>International Workshop on Texture Analysis and Synthesis</i> , pp. 59-63, October 2003.
Kreyszig1983	E. Kreyszig. Advanced Engineering Mathematics. 5 th Edition. Wiley, 1983.
Kube1988	P. Kube and A. Pentland. On the imaging of fractal surfaces. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 10(5), pp.704-707, September 1988.
Lafortune1997	E.P.F. Lafortune, S-C. Foo, K.E. Torrance and D.P. Greenberg. Non-linear approximation of reflectance functions. In <i>Proceedings of the 24th Conference on Computer Graphics and</i> <i>Interactive Techniques</i> , pp. 117-126, 1997.
Lalonde1997	P. Lalonde and A. Fournier. A wavelet representation of reflectance functions. <i>IEEE Transactions on Visualization and Computer Graphics</i> , 3(4), pp. 329-336, 1997.
Lee1993	K.M. Lee and C-C.J. Kuo. Surface reconstruction from photometric stereo images. <i>Journal of the Optical Society of America A</i> , 10(5), pp. 855-868, May 1993.
Leung2001	T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. <i>International Journal of Computer Vision</i> , 43(1), pp. 29-44, 2001.

Lin1999	S. Lin and S.W. Lee. Estimation of diffuse and specular appearance. In <i>Proceedings of the</i> 7 th <i>IEEE International Conference on Computer Vision</i> , 2, pp.20-27, September 1999.
Linnett1991	L.M. Linnett. Multi-texture image segmentation. <i>Ph.D. Thesis</i> , Department of Electrical and Electronic Engineering, Heriot-Watt University, 1991.
Lu1998	R. Lu, J.J. Koenderink and A.M.L. Kappers. Optical properties (bidirectional reflection distribution functions) of velvet. <i>Applied Optics</i> , 37(25), September 1998.
Lu2000	R. Lu, J.J. Koenderink and A.M.L. Kappers. Optical properties (bidirectional reflection distribution function) of shot fabric. <i>Applied Optics</i> , 39(31), November 2000.
Malzbender2001	T. Malzbender, D. Gelb and H. Wolters. Polynomial texture maps. In <i>Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques</i> , pp. 519-528, 2001.
Marschner1999	S.R. Marschner, S.H. Westin, E.P.F. Lafortune, K.E. Torrance and D.P. Greenberg. Image-based BRDF measurement including human skin. In <i>Proceedings of the 10th Workshop on Rendering</i> , <i>Eurographics</i> , pp. 131-144, June 1999.
McGunnigle1998	G. McGunnigle. The classification of texture surfaces under varying illumination direction. <i>Ph.D. thesis</i> , Dept. of Computing and Electrical Engineering, Heriot-Watt University, 1998.
McGunnigle2001	G. McGunnigle and M. J. Chantler. Recovery of Fingerprints using Photometric Stereo. In <i>Proceedings of Irish Machine Vision and Image Processing Conference</i> , pp192-199, September 2001.
McGunnigle2001	G. McGunnigle and M.J. Chantler. Segmentation of rough surfaces using reflectance. In <i>Proceedings of British Machine Vision Conference</i> , pp.323-332, September 2001.
Mulvaney1989	D.J. Mulvaney, D.E. Newland and K.F. Gill. A complete description of surface texture profiles. <i>Wear</i> , 132, pp. 173-182, 1989.
Nayar1990	S.K. Nayar, K. Ikeuchi and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric sampling. <i>IEEE Transactions on Robotics and Automation</i> , 6(4), pp.418-431, August 1990.
Nayar1991	S.K. Nayar, K. Ikeuchi and T. Kanade. Surface reflection: physical and geometrical perspectives. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 13(7), 611-634, July 1991.
Nayar1996	S.K. Nayar, X. Fang and T.E. Boult. Separation of reflection components using color and polarization. <i>International Journal of Computer Vision</i> , 21(3), pp. 163-186, 1996.
Nicodemus1977	F.E. Nicodemus, J.C. Richmond, J.J.Hsai, I.W. Ginsberg and T. Limperis. U.S. Department of Commerce, Geometrical considerations and nomenclature for reflectance. <i>National Bureau of Standards Monograph 160, U.S. Department of Commerce</i> , 1977.
-----------------	--
Nishino2001	K. Nishino, Y. Sato and K. Ikeuchi. Eigen-texture method: appearance compression and synthesis based on a 3D model. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 23(11), pp. 1257-1265, 2001.
Onn1990	R. Onn and A. Bruckstein. Integrability disambiguates surface recovery in two-image photometric stereo. <i>International Journal of Computer Vision</i> , 5(1), pp.105-113, 1990.
Oren1995	M. Oren and S. K. Nayar. Generalization of the Lambertian model and implications for machine vision. <i>International Journal of Computer Vision</i> , 14, pp. 227-251, 1995.
Pentland1982	A. P. Pentland. Finding the illuminant direction. <i>Journal of the Optical Society of America</i> , 72(4), April 1982.
Pentland1984	A.P. Pentland. Fractal-based description of natural scenes. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 6(6), pp. 661-673, November 1984.
Petrou2001	M. Petrou and S. Barsky. Shadows and highlights detection in 4- source colour photometric stereo. In <i>Proceedings of the</i> <i>International Conference on Image Processing</i> , 3, pp. 967 – 970, October 2001.
Phong1975	B.T. Phong. Illumination for computer generated pictures. <i>Communications of the ACM</i> , 18(6), pp. 311-317, June 1975.
Pont2003	S.C. Pont and J.J. Koenderink. Split off-specular reflection and surface scattering from woven materials. <i>Applied Optics</i> , 42(8), March 2003.
Press1988	W.H. Press, B.P. Flannery, S.A. Teukolsky and W.T. Vetterling. Numerical Recipes in C. Cambridge University Press, 1988.
Ramamoorthi2002	R. Ramamoorthi. Analytic PCA construction for theoretical analysis of lighting variability in images of a Lambertian object. IEEE <i>Transactions on Pattern Analysis and Machine Intelligence</i> , 24(10), pp. 1322-1333, 2002.
Ramanath2002	R. Ramanath, W.E. Synder, G.L. Bilbro and W.A. Sander. Domosaicking methods for Bayer color arrays. <i>Journal of Electronic Imaging</i> , 11(3), pp.306-315, July 2002.
Robb2003	M. Robb, A.D. Spence, M.J. Chantler and M. Timmins. Real- Time Per-Pixel Rendering of Bump-mapped Textures Captured using Photometric Stereo. In <i>Proceedings of Vision, Video, and</i> <i>Graphics 2003</i> , pp. 79-87, July 2003.
Rushmeier1997	H. Rushmeier, G. Taubin and A. Gueziec. Applying shape from lighting variation to bump map capture. In <i>Proceedings of Eurographics Rendering Workshop 1997</i> , pp. 35–44, 1997.

Rushmeier1999	H. Rushmeier and F. Bernadini. Computing consistent normals and colors from photometric data. In <i>Proceedings of the</i> 2^{nd} <i>International Conference on 3-D Digital Imaging and Modeling</i> , pp. 99 – 108, October 1999.
Saito1996	H. Saito, K. Omata and S. Ozawa. Recovery of shape and surface reflectance of specular object from rotation of light source. In <i>Proceedings of</i> 2^{nd} <i>International Conference on 3D Digital Imaging and Modeling</i> , pp. 526-535, October 1996
Saltelli2000	A. Saltelli, K. Chan and E.M. Scott. Sensitivity Analysis. Wiley, 2000.
Sato1997	Y. Sato, M.D. Wheeler and K. Ikeuchi. Object shape and reflectance modeling from observation. In <i>Proceedings of 24th Annual Conference of Computer Graphics, SIGGRAPH</i> '97, pp. 379-388, August 1997.
Schlüns1993	K. Schlüns and O. Wittig. Photometric stereo for non- Lambertian surfaces using color information. In <i>Proceedings of</i> <i>the</i> 5 th <i>International Conference on Computer Analysis of Images</i> <i>and Patterns</i> , pp. 444-451, September 1993.
Schlüns1997	K. Schlüns. The irradiance error and its effect in photometric stereo. In <i>Proceedings of the First Joint Australia & New Zealand Biennial Conference on Digital Image & Vision Computing: Techniques and Applications</i> , pp. 539-544, December 1997.
Schlüns1997	K. Schlüns. Shading based 3D shape recovery in the presence of shadows. In <i>Proceedings of the First Joint Australia & New Zealand Biennial Conference on Digital Image & Vision Computing: Techniques and Applications</i> , pp. 195-200, December 1997.
Shafer1985	S.A. Shafer. Using color to separate reflection components. <i>Color Research and Applications</i> , 10(4), pp.210-218, 1985.
Shashua1992	A. Shashua. Geometry and photometry in 3D visual recognition. <i>Ph.D. thesis</i> , Massachusetts Institute of Technology, November 1992.
Smith1999a	M.L. Smith. Gradient space analysis of surface defects using a photometric stereo derived bump map. <i>Image Vision and Computing</i> , 17, pp. 321-332, 1999
Smith1999b	M.L. Smith. The analysis of surface texture using photometric stereo acquisition and gradient space domain mapping. <i>Image Vision and Computing</i> , 17, pp. 1009-1019, 1999
Smith2000	M.L. Smith, A.R. Farooq, L.N. Smith and P.S. Midha. Surface inspection using conventional techniques applied to a photometrically acquired bump map. <i>Sensor Review</i> , 20(4), pp. 299-306, 2000.

Solomon1996	F. Solomon and K. Ikeuchi. Extracting the shape and roughness of specular lobe objects using four light photometric stereo. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 18(4), pp. 449-454, April 1996.
Spence2003a	A. D. Spence and M.J. Chantler. Optimal illumination for three- image photometric stereo acquisition of surface texture. In <i>Proceedings of the</i> 3^{rd} <i>International Workshop on Texture</i> <i>Analysis and Synthesis</i> , pp. 89-94, October 2003.
Spence2003b	A. Spence and M.J. Chantler. On capturing 3D isotropic surface texture using uncalibrated photometric stereo. In <i>Proceedings of the 3rd International Workshop on Texture Analysis and Synthesis</i> , pp. 83-87, October 2003.
Spence2004	A. Spence, M. Robb, M. Timmins and M.J. Chanter. Real-time per-pixel rendering of textiles for virtual textile catalogues. <i>International Journal of Clothing Science and Technology</i> , 16(1), 51-62, 2004.
Tagare1991	H.D. Tagare and R.J.P. deFigueiredo. A theory of photometric stereo for a class of diffuse non-Lambertian surface. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 13(2), February 1991.
Torrance1967	K.E. Torrance and E.M. Sparrow. Theory for off-specular reflection from roughened surfaces. <i>Journal of the Optical Society of America</i> , 57, pp.1105-1114, 1967.
Turk1991	M. Turk and A. Pentland. Eigenface for recognition. <i>Journal of Cognitive Neuroscience</i> , 3(1), pp.71-96, 1991.
Varma2004	M. Varma and A. Zisserman. Estimating illumination direction from textured images. In <i>Proceedings of IEEE Conference on Computer Vision and Pattern Recognition</i> , pp.179-186, 2004.
Ward1992	G.J. Ward. Measuring and modeling anisotropic reflection. In <i>Proceedings of Computer Graphics, SIGGRAPH</i> '92, 26(2), pp. 255-263, 1992.
Westin1992	S.H. Westin, J.R. Arvo and K.E. Torrance. Predicting reflectance functions from complex surfaces. Computer Graphics, 26(2), pp. 255-264, July 1992
Wolff1994	L.B. Wolff. Diffuse reflectance model for smooth dielectric surfaces. <i>Journal of the Optical Society of America A</i> , 11(11), pp. 2956-2968, November 1994.
Wolff1998	L.B. Wolff, S.K. Nayar and M. Oren. Improved diffuse reflection models for computer vision. <i>International Journal of Computer Vision</i> , 30(1), pp. 55-71, 1998.
Woo1999	M. Woo, J. Neider, T. Davis and D. Shreiner. OpenGL programming guide: the official guide to learning OpenGL, Version 1.2, 3 rd Edition, Addison Wesley Professional, 1999.

Woodham1980	R. J. Woodham. Photometric method for determining surface orientation from multiple images. <i>Optical Engineering</i> . 19(1), pp. 139-144, 1980.
Woodham1991	R.J. Woodham, Y. Iwahori and R.A. Barman. Photometric stereo: Lambertian reflectance and light sources with unknown direction and strength. <i>Technical Report 91-18</i> , Laboratory for Computational Intelligence, University of British Columbia, Vancouver, Canada, August 1991.
Woodham1994	R. J. Woodham. Gradient and curvature from the photometric- stereo method, including local confidence estimation. <i>Journal of</i> <i>the Optical Society of America A</i> , 11(11), pp. 3050-3068, November 1994.
Yuille1997	A. Yuille and D. Snow. Shape and albedo from multiple images using integrability. In <i>IEEE Proceedings of Computer Vision and Pattern Recognition</i> , pp. 158-164, 1997.
Yuille1999	A. L. Yuille, D. Snow, R. Epstein and P. N. Belhumeur. Determining generative models of objects under varying illumination : shape and albedo from multiple images using SVD and integrability. <i>International Journal of Computer Vision</i> , 35(3), pp. 203–22, 1999.
Zhang1997	R. Zhang and M. Shah. Iterative shape recovery from multiple images. <i>Image and Vision Computing</i> , 15, pp. 801-81, 1997.
Zhang1999	R. Zhang, P-S. Tsai, J.E. Cryer and M. Shah. Shape from shading: a survey. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 21(8), August 1999.
Zheng1991	Q. Zheng and R. Chellapa. Estimation of illuminant direction, albedo and shape from shading. <i>IEEE Transactions on Pattern Analysis and Machine Intelligence</i> , 13, pp. 915-920, 1991.