

# Height Recovery of Rough Surfaces from Intensity Images

*Cristina Gullón*

Doctor of Philosophy

Heriot-Watt University  
School of Engineering and Physical Sciences

February, 2003

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 written consent of the author or the University (as may be appropriate).

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Motivation . . . . .	1
1.2	Scope . . . . .	2
1.3	Contribution . . . . .	4
1.4	Thesis organisation . . . . .	5
<b>2</b>	<b>Topographic Texture</b>	<b>8</b>
2.1	Definitions . . . . .	9
2.2	Description of textures . . . . .	13
2.3	Synthetic textures . . . . .	14
2.4	Real textures . . . . .	19
2.4.1	Description . . . . .	19
2.4.2	Do the real textures comply with the surface assumptions? . .	23
2.4.3	Texture sampling . . . . .	27
2.5	Discussion . . . . .	29
<b>3</b>	<b>Surface-to-Image Models</b>	<b>31</b>
3.1	Definitions . . . . .	32
3.2	Theory of light reflection . . . . .	34
3.3	Lambert's law . . . . .	36
3.4	Modelling specularities: Phong's . . . . .	38
3.5	Physically based models . . . . .	41
3.6	Modelling shadows . . . . .	42
3.7	Kube's model . . . . .	44
3.8	Assessment of reflectance models . . . . .	45
3.8.1	Validity of Kube's model . . . . .	46
3.8.2	What is the effect of specularities? . . . . .	48
3.9	Discussion . . . . .	51

<b>4</b>	<b>Shape Recovery: A Review</b>	<b>52</b>
4.1	Shape from Shading . . . . .	53
4.1.1	Introduction . . . . .	53
4.1.2	Review of single-light Shape from Shading techniques . . . . .	54
4.2	Photometric Stereo . . . . .	57
4.2.1	Introduction . . . . .	57
4.2.2	Review of Photometric Stereo techniques . . . . .	57
4.3	Integration techniques . . . . .	60
4.4	Discussion . . . . .	63
4.5	Benchmark photometric recovery technique . . . . .	63
4.5.1	Simple Photometric Stereo scheme . . . . .	64
4.5.2	Frequency integration . . . . .	65
4.6	Summary and discussion . . . . .	67
<b>5</b>	<b>Shape Recovery: Two Novel Methods</b>	<b>68</b>
5.1	Two-light Photometric Stereo techniques . . . . .	69
5.1.1	Review of the algorithms . . . . .	69
5.1.2	Uniqueness and existence . . . . .	71
5.2	Linear Photometric Stereo . . . . .	73
5.2.1	Introduction . . . . .	73
5.2.2	Description of the algorithm . . . . .	74
5.2.3	Advantages . . . . .	80
5.3	Optimal Linear Filter . . . . .	80
5.3.1	Introduction . . . . .	80
5.3.2	Description of the algorithm . . . . .	82
5.3.3	Advantages . . . . .	86
5.4	Assumptions . . . . .	87
5.5	Summary and discussion . . . . .	88
<b>6</b>	<b>Optimal Lighting Conditions</b>	<b>90</b>
6.1	Chapter organisation . . . . .	91
6.2	Optimal light source position . . . . .	91
6.2.1	What is the optimal position for two lights? . . . . .	91
6.2.2	What are the best azimuths for $n$ lights? . . . . .	100
6.3	Trade-off on the number of lights . . . . .	103
6.3.1	Introduction . . . . .	103
6.3.2	Simulation . . . . .	103
6.4	Summary and discussion . . . . .	104

<b>7</b>	<b>Assessment of Surface Estimation: Simulation</b>	<b>106</b>
7.1	Simulation Framework . . . . .	107
7.2	Effect of Surface Roughness . . . . .	108
7.2.1	Motivation . . . . .	108
7.2.2	Method . . . . .	109
7.2.3	Results and discussion . . . . .	110
7.3	Robustness to shadowing . . . . .	111
7.3.1	Motivation . . . . .	111
7.3.2	Method . . . . .	111
7.3.3	Results and discussion . . . . .	112
7.4	Effect of surface reflectance . . . . .	114
7.4.1	Motivation . . . . .	114
7.4.2	Method . . . . .	115
7.4.3	Results and discussion . . . . .	116
7.5	Robustness against noise . . . . .	119
7.5.1	Motivation . . . . .	119
7.5.2	Method . . . . .	120
7.5.3	Results and discussion . . . . .	121
7.6	Discussion . . . . .	121
<b>8</b>	<b>Assessment of Surface Estimation: Experiment</b>	<b>124</b>
8.1	Assessment of the recovery through rendering . . . . .	125
8.1.1	Experimental method . . . . .	125
8.1.2	Experimental results . . . . .	129
8.2	Surface estimation for image prediction purposes . . . . .	138
8.2.1	Experimental method . . . . .	138
8.2.2	Results and discussion . . . . .	140
8.3	Summary and discussion . . . . .	145
<b>9</b>	<b>Improving the Surface Reconstruction</b>	<b>148</b>
9.1	Chapter organisation . . . . .	149
9.2	How could the Photometric estimate be improved? . . . . .	149
9.2.1	Regularisation schemes: literature review . . . . .	149
9.2.2	Discussion . . . . .	151
9.3	A simple iterative method . . . . .	152
9.3.1	Description of the algorithm . . . . .	152
9.3.2	When to stop iterating? . . . . .	154
9.3.3	Summary . . . . .	155
9.4	Assessment of the iterative recovery through rendering . . . . .	155



9.4.1	Experimental method . . . . .	155
9.4.2	Experimental results . . . . .	157
9.5	Conclusions . . . . .	165
<b>10</b>	<b>Conclusions</b>	<b>166</b>
10.1	Summary . . . . .	166
10.2	Conclusions . . . . .	169
<b>A</b>	<b>Slope distribution for the real surfaces</b>	<b>171</b>
<b>B</b>	<b>Result tables</b>	<b>179</b>
B.1	Should we use albedo for rendering the Linear estimate? . . . . .	179
B.2	Assessment of surface estimation through rendering . . . . .	181
B.3	Reduction of the brightness error with the iteration process . . . . .	183
B.4	Accuracy of prediction with the iteration process . . . . .	183
B.5	Assessment of the iterative Linear estimate through rendering . . . . .	186
<b>C</b>	<b>Reconstruction results</b>	<b>188</b>
	<b>References</b>	<b>207</b>

# List of Figures

2.1	Synthetic surfaces: Sphere (left) and Rock (right)	11
2.2	(a) Surface profiles; (b) Surface of Ganymede	12
2.3	Synthetic textures: Fractal (left), Mulvaney (centre) and Ogilvy (right)	17
2.4	Polar plot of two Ogilvy surfaces	18
2.5	Real textures	20
2.6	Real textures	22
2.7	Slope distributions of real surfaces (real and fitted Gaussian)	28
3.1	The geometry of the imaging system	32
3.2	Masking (a) and self and cast shadows (b)	33
3.3	(a) Solid angle; (b) Foreshortened area	34
3.4	A BDRF relates incoming energy ( $\phi_i, \theta_i$ ) to outgoing intensity ( $\phi_v, \theta_v$ )	34
3.5	Reflection: (a) mirror-like specular (b) specular+diffuse (c) perfectly diffuse	36
3.6	Mirror direction for Phong's model	40
3.7	Example of self and cast shadows in a sinusoid	43
3.8	Validity of Kube's approximation to Lambert's law	48
3.9	Phong's specular term contribution to total reflection against $K_d$	49
3.10	Phong's specular term contribution to total reflection against $n$	50
5.1	Reflectance maps of a Lambertian surface	76
5.2	Polar plots of the spectra of two images of an isotropic fractal obtained rendering with Lambert's law and Kube's model.	77
5.3	Surface recovery from 5 intensity images	79
5.4	Magnitude of the frequency response for the ideal derivative operator and the central difference operator	83
6.1	Contour plots of Lambertian reflectance maps	94
6.2	LPS estimation accuracy against the position of the lights	97
6.3	OLF estimation accuracy against the position of the lights	97

6.4	LPS estimation accuracy against the position of the lights relative to the surface grain . . . . .	99
6.5	OLF estimation accuracy against the position of the lights relative to the surface grain . . . . .	99
6.6	Dependency of intensity with $(\theta - \tau)$ . . . . .	102
7.1	General scheme for the surface estimation assessment . . . . .	108
7.2	Procedure for assessing surface estimation against surface roughness .	109
7.3	Surface recovery performance against rms slope . . . . .	110
7.4	Procedure for assessing surface estimation against shadowing . . . . .	112
7.5	3PS performance against shadowing . . . . .	112
7.6	LPS performance against shadowing . . . . .	113
7.7	OLF performance against shadowing . . . . .	114
7.8	Procedure for assessing surface estimation against surface reflectance	116
7.9	Recovery performance against percentage of diffuse reflection . . . . .	118
7.10	Recovery performance against specular peak shape . . . . .	119
7.11	Fractal (left), Mulvaney (centre) and Ogilvy (right). Top row: no noise. Second row: $S/N = 10dB$ . Bottom row: $S/N = 0dB$ . . . . .	120
7.12	Procedure for assessing surface estimation against noise . . . . .	121
7.13	Recovery performance for synthetic textures against noise . . . . .	122
8.1	Recovery accuracy of the Linear PS and the 3-light PS techniques . .	130
8.2	<i>Moderate roughness fracture</i> imaged for $\sigma = 45^\circ$ and $\tau = 300^\circ$ . . . .	131
8.3	Estimated <i>moderate roughness fracture</i> surface . . . . .	131
8.4	<i>Rough fracture</i> imaged for $\sigma = 45^\circ$ and $\tau = 300^\circ$ . . . . .	133
8.5	Estimated <i>rough fracture</i> surface . . . . .	133
8.6	Estimated <i>heavy deposited texture</i> . . . . .	135
8.7	<i>Repetitive primitive texture</i> imaged for $\sigma = 45^\circ$ and $\tau = 330^\circ$ . . . . .	137
8.8	Estimated <i>repetitive primitive texture</i> . . . . .	138
8.9	<i>Heavy deposited surface</i> : rendering example . . . . .	141
8.10	<i>Rough deposited fracture</i> : Accuracy of prediction when rendering the 3-light Photometric Stereo estimate . . . . .	142
8.11	<i>Rough deposited fracture</i> : Accuracy of prediction when rendering the Linear Photometric Stereo estimate . . . . .	143
8.12	<i>Sand-rippled surface</i> : rendering example . . . . .	144
8.13	<i>Sand-rippled surface</i> : Accuracy of prediction when rendering the Linear Photometric Stereo estimate . . . . .	145
8.14	<i>Sand-rippled surface</i> : Accuracy of prediction when rendering the 3-light Photometric Stereo estimate . . . . .	146

9.1	Accuracy of image prediction for the 3-light PS and the Iterative 3-light PS techniques . . . . .	158
9.2	Accuracy of image prediction for the Linear PS and the Iterative Linear PS techniques . . . . .	158
9.3	Accuracy of recovery for the Iterative Linear and the 3-light PS techniques . . . . .	159
9.4	<i>Gentle roughness fracture</i> imaged for $\sigma = 45^\circ$ and $\tau = 300^\circ$ . . . . .	160
9.5	Estimated <i>gentle roughness fracture</i> surface . . . . .	161
9.6	<i>Sparsely deposited texture</i> imaged for $\sigma = 45^\circ$ and $\tau = 300^\circ$ . . . . .	162
9.7	Estimated <i>sparsely deposited texture</i> . . . . .	163
9.8	Estimated <i>high frequency sand ripple</i> surface . . . . .	164
A.1	Slope distributions of real surfaces . . . . .	173
A.2	Slope distributions of real surfaces . . . . .	174
A.3	Slope distributions of real surfaces . . . . .	175
A.4	Slope distributions of real surfaces . . . . .	176
A.5	Slope distributions of real surfaces . . . . .	177
A.6	Slope distributions of real surfaces . . . . .	178
B.1	Accuracy of recovery for the Iterative 3-light PS and the Iterative Linear PS techniques . . . . .	186
C.1	<i>Rough fracture</i> : rendering example . . . . .	190
C.2	Height recovery for the <i>rough fracture</i> . . . . .	190
C.3	<i>Moderate roughness fracture</i> : rendering example . . . . .	191
C.4	Height recovery for the <i>Moderate roughness fracture</i> . . . . .	191
C.5	<i>Gentle roughness fracture</i> : rendering example . . . . .	192
C.6	Height recovery for the <i>gentle roughness fracture</i> . . . . .	192
C.7	<i>Smooth cured fracture</i> : rendering example . . . . .	193
C.8	Height recovery for the <i>smooth cured fracture</i> . . . . .	193
C.9	<i>Smooth patterned fracture</i> : rendering example . . . . .	194
C.10	Height recovery for the <i>smooth patterned fracture</i> . . . . .	194
C.11	<i>Moderate roughness deposited fracture</i> : rendering example . . . . .	195
C.12	Height recovery for the <i>moderate roughness deposited fracture</i> . . . . .	195
C.13	<i>Rough deposited fracture</i> : rendering example . . . . .	196
C.14	Height recovery for the <i>Rough deposited fracture</i> . . . . .	196
C.15	<i>Heavily deposited fracture</i> : rendering example . . . . .	197
C.16	Height recovery for the <i>heavily deposited fracture</i> . . . . .	197
C.17	<i>Medium deposited fracture</i> : rendering example . . . . .	198

C.18 Height recovery for the <i>medium deposited fracture</i> . . . . .	198
C.19 <i>Light deposited fracture</i> : rendering example . . . . .	199
C.20 Height recovery for the <i>light deposited fracture</i> . . . . .	199
C.21 <i>Sparse deposited fracture</i> : rendering example . . . . .	200
C.22 Height recovery for the <i>sparse deposited fracture</i> . . . . .	200
C.23 <i>Sand-rippled surface</i> : rendering example . . . . .	201
C.24 Height recovery for the <i>sand-rippled surface</i> . . . . .	201
C.25 <i>Sand-rippled surface (high frequency)</i> : rendering example . . . . .	202
C.26 Height recovery for the <i>sand-rippled surface (high frequency)</i> . . . . .	202
C.27 <i>Smoothed sand-rippled surface</i> : rendering example . . . . .	203
C.28 Height recovery for the <i>smoothed sand-rippled surface</i> . . . . .	203
C.29 <i>Net anaglypta</i> : rendering example . . . . .	204
C.30 Height recovery for the <i>net anaglypta</i> . . . . .	204
C.31 <i>Rippled anaglypta</i> : rendering example . . . . .	205
C.32 Height recovery for the <i>rippled anaglypta</i> . . . . .	205
C.33 <i>Rice grain shape anaglypta</i> : rendering example . . . . .	206
C.34 Height recovery for the <i>rice grain shape anaglypta</i> . . . . .	206

# List of Tables

2.1	Roughness measures for the smooth sphere and the rough fractal . . .	12
2.2	Description of 6 Ogilvy textures that comply with the unidirection- ality condition . . . . .	18
2.3	Gradient mean and roughness for the real surfaces . . . . .	25
2.4	Validity of the integrability assumption for the real textures . . . . .	26
2.5	Photometric sampling for real textures . . . . .	29
5.1	Two-light Photometric Stereo techniques . . . . .	69
5.2	Summary of the assumptions made during the derivation of the three considered techniques . . . . .	88
6.1	Recovery accuracy against number of images . . . . .	104
8.1	Pixel-width of the image of the real textures [mm] . . . . .	129
8.2	Average accuracy of recovery for fractures [dB] . . . . .	132
8.3	Average accuracy of recovery for deposited surfaces [dB] . . . . .	134
8.4	Average accuracy of recovery for sand-rippled and anaglypta surfaces [dB] . . . . .	136
8.5	Average accuracy of recovery for special cases [dB] . . . . .	137
8.6	Accuracy of the prediction of the surface appearance for isotropic surfaces [dB] . . . . .	140
8.7	Accuracy of the prediction of the surface appearance for directional surfaces [dB] . . . . .	143
9.1	Average accuracy of recovery for fractures [dB] . . . . .	160
9.2	Average accuracy of recovery for deposited surfaces [dB] . . . . .	161
9.3	Average accuracy of recovery for sand-rippled and anaglypta surfaces [dB] . . . . .	163
9.4	Average accuracy of recovery for special textures [dB] . . . . .	165
B.1	Accuracy of the predicting the real surface appearance using the Lin- ear estimate [dB] . . . . .	180

B.2	Accuracy of the prediction of real surfaces appearance [dB] . . . . .	182
B.3	Brightness error before and after iteration for the 3-light PS and the Linear PS [dB] . . . . .	184
B.4	Accuracy of prediction before and after iteration for the Linear PS and the 3-light PS [dB] . . . . .	185
B.5	Accuracy of prediction of the Iterative Linear PS technique against zenith . . . . .	187

# List of Symbols

$\delta$ or $R_q$	Rms roughness
$R_a$	Average roughness
$p_{rms}$	Rms slope in the $x$ direction
$q_{rms}$	Rms slope in the $y$ direction
$d$	Surface directionality
$s$	Surface
<b>S</b>	Surface gradient field
$S$	Surface power spectrum
$\omega$	Radial frequency
$\theta$	Polar frequency angle
$u$	Horizontal frequency index
$v$	Vertical frequency index
$\beta$	Surface power roll-off
$\omega_c$	Cut-off frequency
$\lambda_1$	Correlation length in the $x$ direction
$\lambda_2$	Correlation length in the $y$ direction
<b>N</b>	Surface normal unit vector
<b>L</b>	Illuminant direction unit vector
<b>V</b>	Viewing direction unit vector
<b>R</b>	Mirror direction unit vector
$\sigma$	Zenith angle of the light source
$\tau$	Azimuth angle of the light source
$(k_a, k_b, k_c)$	Light source direction
<b>I</b> or $i$	Reflected intensity
$I_0$ or $i_0$	Incident intensity
$\rho$	Albedo
$p$	Facet slope in the $x$ direction
$q$	Facet slope in the $y$ direction
$n$	Phong's parameter to model roughness
$K_d$	Diffuse reflection coefficient



$K_s$	Specular reflection coefficient
$R(p, q)$	Reflectance map
$E(x, y)$	Image
$a, b, c$	Parameters of the Optimal Linear Filter
$i_d$	Desired image
$e_{rms}$	Rms error
$e$	Error function
$e_s$	Error due to departure from smoothness
$e_i$	Brightness error
$\alpha$	Regularisation parameter

# Acknowledgements

I would like to thank my supervisor, Dr. Mike Chantler, for his guidance and advice throughout this thesis. His support, availability, enthusiasm and long discussions have been invaluable.

My thanks are also due to Dr. Laurie Linnett, Dr. Mike Chantler and Dr. Gerald McGunnigle for making their earlier work, software and data available to me.

I would also like to express my gratitude to Ged, Mike, Ivan, Andy and Dong for proof-reading my thesis.

To the folks in the ISL and the texture lab, thanks for teaching me what research is all about.

I am especially indebted to Dr. Gerald McGunnigle for all his comments, suggestions and interesting discussions during this research.

Nunca habria considerado este proyecto sin vuestro apoyo. ¡Gracias! ¡Os quiero!

# Abstract

This thesis is concerned with the 3D estimation of rough surfaces from their intensity images. A technique which combines Photometric Stereo and frequency integration is proposed. The combination of these two standard methods for reconstructing rough surfaces is novel. We refer to this technique as the Benchmark technique. Two novel recovery algorithms which rely on assumptions about the linearity of the surface reflectance are also presented. We refer to them as the Optimal Linear Filter and the Linear Photometric Stereo. The proposed methods differ in the information that they require as well as in the assumptions that they make about the surface.

The ability of the proposed techniques to estimate rough surfaces is assessed using simulation and real data. The assessment considers a diverse set of textures including those that are challenging for the algorithms, such as very rough or specular surfaces. The most robust estimation is given by the Optimal Linear Filter. However this technique requires information about the surface topography, which is usually not available. Between the alternatives, the Benchmark technique gives more accurate reconstructions.

A post-processing step which can be used to improve the surface estimate is presented. This minimises the brightness error using an iterative approach. When the Linear Photometric Stereo method is combined with the post-processing step, its performance is similar to that of the Benchmark technique, despite requiring one less image. However the Linear Photometric Stereo algorithm is restricted to constant albedo surfaces. The choice of the most appropriate method is determined by the application requirements.

# Chapter 1

## Introduction

---

### 1.1 Motivation

This thesis is concerned with the non-contact recovery of rough test surfaces. Our aim is to reconstruct an accurate three-dimensional heightmap of the surface. There are many applications of surface recovery, e.g. surface inspection [1], corrosion monitoring [2], fault detection [3], face recognition [4], segmentation [5] and classification [6], automated assembly [7].

Several non-contact methods have been proposed for measuring depth [8]. They can be divided into direct and indirect methods. *Direct methods* measure range explicitly [9]. Time-of-flight techniques [10], for instance, measure range by estimating how long it takes light to reach the target and return to the sensor. A short, intense pulse, emitted from a laser, travels toward the surface, and a fraction of the signal is scattered back to the sensor. This approach is known as pulsed laser time-of-flight. Phase-shift based systems send a modulated optical signal to the surface. Range is measured by calculating the phase shift of the received signal. Depth information is the only information available in time-of-flight systems. They are usually complicated and expensive systems that take a long time to measure a 3D surface shape accurately.

*Indirect methods* estimate distance by measuring parameters calculated from intensity images of the illuminated object. Information about the reflectance or the lighting conditions of the object can also be determined with these systems. We discuss two categories of indirect approaches: Shape from Shading techniques infer

the shape of surface facets [11] from their intensity; the second category estimates range by measuring how two or more images correspond in terms of focus and spatial position. Stereo Vision, Optical Flow, Shape from Texture and Shape from Focus belong to this second category.

*Stereo Vision* uses triangulation to compute depth [12]. *Optical flow* calculates the relative distance to points on the surface of an object by analysing how image points flow from one frame to the next [13]. The problem with these two algorithms is that they require a relationship between image points in one frame and the same points in another frame. This is known as the correspondence problem. The objective of *Shape from Texture* is to estimate the orientation of a scene plane from a perspective image under the assumption that the scene is coated with texture [14] [15]. *Shape from Focus* methods use a sequence of images taken by a single camera at different focus levels to compute depth of objects [16] [17]. These two techniques are more suitable to recover the shape of objects than the structure of a rough surface.

*Shape from Shading* techniques measure the surface shape by analysing the radiometry of image formation [18]. This method requires some control over the lighting conditions of the scene. However, it is a simple, computationally inexpensive, fast technique that permits an accurate reconstruction of the surface shape. Its operating range varies from a few millimetres to thousands of meters depending on the application. It is this category that is the subject of this thesis.

## 1.2 Scope

This work treats the problem of recovering the three-dimensional structure of a rough, textured, surface. Among the reviewed non-contact methods, we found Shape from Shading techniques most suitable. The problem becomes one of estimating the surface shape from its images.

Although Shape from Shading has been extensively used to recover the 3D struc-

ture of a scene, it is been mostly applied to smooth surfaces [18] [19]. In this thesis, we investigate the less thoroughly researched problem of recovering rough surfaces. We estimate depth and not gradient. Neither do we seek to recover the reflectance of the surface nor the lighting conditions.

The observed image is a function of the surface. However, it is also dependent on other variables. In the case of rough textured surfaces, the observed image is the *perceived texture*. This consists of intensity variations in the image plane which are due to a projected illumination pattern, surface markings<sup>1</sup>, the interaction of illumination with a rough surface or any combination of these. Furthermore, the surface image depends on the geometry of the imaging system and on the imaging device. In defining the scope of this research it is useful to describe the assumptions on the imaging process as well as on the characteristics of the test surfaces:

1. *The camera* is static throughout the imaging process and it is placed directly above the surface. A Vöskühler<sup>2</sup>, 12 bit monochrome camera with constant gain is used.
2. The projection of the surface onto the CCD array is assumed to be orthographic.
3. *The source* is a point light source at infinity, so that the illumination incident on the scene is considered to be uniform in direction and intensity.
4. The position of the light is controlled so that the surface can be lit from different directions.
5. *The surface* is globally flat, i.e. its local variation is limited and each of its partial derivative fields sums to approximately zero.
6. Furthermore, it is continuous and rough with low slope angles. The surface roughness is due to a lack of correlation amongst neighbouring slopes.

---

<sup>1</sup>We refer to the surface marking as the *albedo texture*.

<sup>2</sup>[www.vdsvosk.de/en/](http://www.vdsvosk.de/en/)

7. The surface is assumed to have a uniformly matte or Lambertian [20] reflectance, i.e. it contains only *topographic texture*<sup>3</sup>.

The *perceived texture* is due only to surface relief. Thus we refer to *texture* meaning surface topography. This is a consequence of assuming a uniform intensity (point 3) and a constant albedo surface (point 7).

## 1.3 Contribution

The author believes that the contributions of this thesis are fourfold:

- In this work a problem that has been neglected in the literature of Shape from Shading is tackled: 3D estimation applied to rough surfaces. The assumptions about the surface smoothness adopted by most Shape from Shading techniques, make them unsuitable for rough surfaces; though some of these algorithms have been tested on rough surfaces and have shown an acceptable performance. However it is the author's belief that an approach which does not rely on smoothness assumptions provides a better choice for rough surface reconstruction.
- The combination of Photometric Stereo estimation with global frequency integration to recover rough surfaces.
- Two novel recovery algorithms are proposed. Both are simple techniques which require few intensity data to solve for surface height. And, importantly, both can be applied to rough surfaces.
- Finally, an assessment of the novel recovery algorithms is provided. Their performance is compared to that of the Benchmark technique on synthetic and real data.

---

<sup>3</sup>Term used to refer to the three-dimensional variation, or relief of a surface.

## 1.4 Thesis organisation

**Chapter two** is concerned with the topographic shape of the test surfaces. The chapter starts by defining relevant terms, such as texture, surface roughness and surface directionality. The synthetic and real test textures are then introduced. The basic characteristics and assumptions about the textures are highlighted. Three well-known synthetic surface models are defined. These models are parametric and can be used to generate a variety of rough surfaces. A description of the real data which covers their topographies, reflectance characteristics and degrees of directionality and roughness is introduced. The validity of the surface assumptions on our real data is investigated.

**Chapter three** reviews surface-to-image models. First a definition of shadowing is given, some radiometric terms are defined and the imaging system is described. Then the background theory of light reflection is reviewed. Existing reflection models are surveyed. The three most appropriate models for our simulations are more extensively described. Finally, the described reflection models are assessed through the use of simulation.

**Chapter four** surveys the existing literature of Shape from Shading techniques. Special attention is paid to one of these methods: Photometric Stereo. The advantages of using this method for the recovery of rough surface are discussed. Since the aim of this thesis is height estimation, the surface gradient has to be integrated. Integration techniques are then reviewed concluding that global integration is more appropriate for our purposes. In this chapter a benchmark recovery technique is described. It combines a successful Photometric Stereo technique with a well-known frequency integration algorithm. The combination of these two standard components make the recovery technique suitable for rough surfaces.

**Chapter five** introduces our contribution to shape recovery. The chapter starts



with a survey into 2-light Photometric Stereo techniques. The suitability of the reviewed techniques for textured surface estimation is discussed. Next two novel two-light shape estimators are proposed: the Linear Photometric Stereo technique and the Optimal Linear Filter. Their advantages for height recovery are reviewed. Several surface assumptions made for the derivation of our algorithms are summarised.

**Chapter six** argues the importance of optimising the lighting conditions to improve surface estimation. The aim of this chapter is to find the best lighting geometry. The optimal number of light sources as well as the best position for them is investigated in this chapter. The conclusions drawn from this investigation are justified both by reviewing the literature and through the use of simulation. An optimised implementation of the novel recovery techniques is then proposed.

**Chapter seven** presents an assessment through simulation of our novel surface estimators. The assessment compares the performance of the proposed recovery methods to that of the Benchmark technique. The robustness of our algorithms is tested against the effect of increasing the surface roughness, varying the surface reflectance, shadowing and noise. The scope of the algorithms and their expected accuracy is then stated. Some limitations of the simulation are reviewed and the necessity for a more realistic assessment is highlighted.

**Chapter eight** presents an assessment of our novel surface estimators on real data. The Linear Photometric Stereo technique is tested and compared to the Benchmark technique. Rendering is presented as a possible approach for assessing the recovery in the absence of fully calibrated data. Surface recovery is assessed against surface roughness and surface discontinuities. Furthermore, a range of textures with different topographies and reflectance is recovered. Some surfaces that severely violate the surface assumptions are also estimated to test the limits of our recovery methods. We argue that the Linear Photometric

Stereo technique is not sufficiently accurate.

**Chapter nine** reviews regularisation techniques to identify an efficient scheme which improves the Photometric estimate. A simple optimisation technique to be used with the Linear Photometric Stereo algorithm is presented. The assessment through rendering of the optimised surface estimates is repeated. The issues of robustness against surface roughness, discontinuities, surface topography and reflectance are again investigated.

**Chapter ten** concludes the thesis. The contributions of this research are summarised and the advantages and limitations of the developed methods are pointed out.

# Chapter 2

## Topographic Texture

---

The algorithms proposed in this thesis use image intensity to estimate the surface topography. The intensity is a function of the surface slope, its reflectance and the illumination conditions. We benefit from the effect that varying the illumination has on the image, to emphasise the surface shape. Furthermore, we want our algorithms to be robust to the type of surface reflectance. This chapter is concerned with the shape of our test surfaces and next chapter explores their reflectance and reviews reflectance models.

This chapter is an introduction to the synthetic surface models and the real textures. We start by defining relevant terms. The basic characteristics and assumptions about our textures are highlighted. We review three synthetic surface models that are well-known in the literature. These synthetic models are used to generate data for our simulations. We describe the real surfaces in terms of their topographies, reflectance characteristics and degrees of directionality and roughness. The validity of the surface assumptions is investigated for our real data.

## 2.1 Definitions

We define the concepts that are important for understanding the problem: we give a working definition of texture; we describe several measures of surface roughness and a measure for directionality; we then give examples of smooth and rough surfaces.

**Texture** is the term we use to characterise the structure and detail of the surface of a given object. It is an important feature for image processing and pattern recognition. Texture can be seen in images ranging from satellite data to microscopic images. Despite its importance and ubiquity in image data, texture lacks a precise definition. One of the most common definitions describe texture as being generated by one or more basic local patterns that are repeated in a periodic manner over some image region. This definition is most applicable to deterministic textures. Images formed by natural physical processes, such as rocks or aerial photographs, rarely possess a basic pattern nor a dominant repetition frequency. However, they are not completely random, although they can be considered stochastic processes. These are the type of textures that are the scope of this thesis.

**Surface roughness** Our textures are globally flat rough surfaces with low slope angles. Thus surface roughness is not due to the presence of high slope facets but to a lack of correlation amongst neighbouring facets. A single parameter is not a sufficient description of surface texture, though surface roughness does measure a characteristic which is relevant. The most common surface statistical parameter is the standard deviation of the heights of a surface, i.e. root-mean-square roughness (rms roughness). For a single line of profile data, it can be obtained as follows:

$$\delta = \sqrt{\frac{1}{n} \sum_{x=1}^n [s(x) - \overline{s(x)}]^2} \quad (2.1)$$

where  $s(x)$  represents the height of a surface at a point  $x$  along the profile,  $\overline{s(x)}$  is the expectation of  $s(x)$  and  $n$  is the number of discrete, equally spaced, measured points along the profile [21]. Rms roughness is also designated by  $R_q$ .

Another common parameter to measure roughness is the average roughness  $R_a$ . It is simply the average of the absolute values of the surface height variations measured from the mean level along a line of profile data [21]. Expressed in equation form, this is:

$$R_a = \frac{1}{n} \sum_{x=1}^n |s(x) - \overline{s(x)}| \quad (2.2)$$

We compute rms roughness and average roughness for all points of the 2D texture. Thus these parameters do not provide information about the surface directionality. We need a parameter to measure directionality. For a continuous surface, the surface slope in a particular direction can be defined as the tangent to the surface in that direction, and the rms slope is a unique well-defined quantity [21]. The directionality of the slope distribution may be parameterised computing the rms slope of the surface in the direction of the  $x$  and  $y$  axes,  $p_{rms}$  and  $q_{rms}$ :

$$p_{rms} = \sqrt{\frac{1}{n} \sum_{x=1}^n \left[ \frac{\partial s(x, y)}{\partial x} - \overline{\frac{\partial s(x, y)}{\partial x}} \right]^2} \quad (2.3)$$

$$q_{rms} = \sqrt{\frac{1}{n} \sum_{y=1}^n \left[ \frac{\partial s(x, y)}{\partial y} - \overline{\frac{\partial s(x, y)}{\partial y}} \right]^2} \quad (2.4)$$

Parameters determined from rough surface measurements can often be subject to artifacts of the measuring process: the measured profile can be affected by measurement variables such as sampling frequency, sample length, discretisation interval, instrument resolutions and sample shape and size. The value of rms slope is particularly dependent on the measuring process [21]. Rms slope, rms roughness and average roughness are only defined for surfaces whose power spectra are band-limited.

**Directionality** We define a single quantity that combines the directional information contained in the vertical and horizontal profiles. We assume that the surface grain is aligned with one of the axes and we measure the ratio:




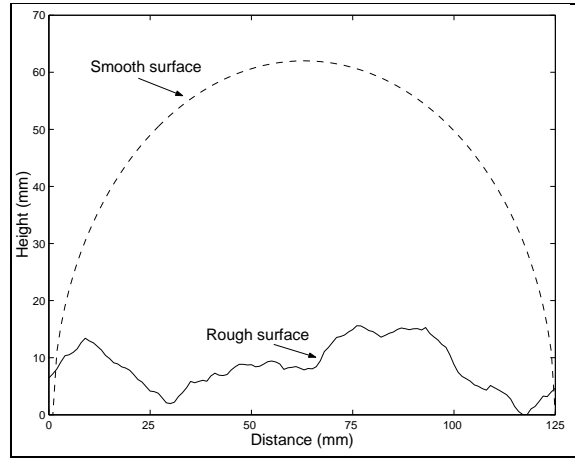
Figure 2.1: Synthetic surfaces: Sphere (left) and Rock (right)

$$directionality = d = \frac{p_{rms}}{p_{rms} + q_{rms}} \quad (2.5)$$


**Smooth and rough surfaces** We define and give examples of smooth and rough surfaces. No measure of roughness can completely describe the surface topography.

Figure 2.1 shows two synthetic surfaces which have the same reflectance characteristics and are imaged under the same lighting. On the left, a hemisphere is rendered: its surface is smooth. The surface on the right is a globally flat rough rock. Figure 2.2 (a) shows a line profile of the surfaces. The roughness measures for both surfaces are shown in Table 2.1. The sphere has higher values of roughness than the rock. The average roughness and the rms roughness are larger for the sphere since it presents large height deviation from the mean level. The rms slope is larger for the sphere because the facets present a wider range of slopes. We only consider globally flat surfaces to avoid this ambiguity. The surface roughness, that we refer to in this thesis, is not due to high slopes but to neighbouring facets with very different low slopes (rapid changes).

**Resolution** We should also bear in mind that augmenting the resolution of the observed surface makes it appear rougher. Figure 2.2 (b) shows the impact craters on the surface of Jupiter’s moon Ganymede. The image was taken by NASA’s Galileo spacecraft. North is to the top of the image and the Sun illuminates the surface from the West. The higher resolution images show that even smooth-looking terrain appears rough at fine scale.



(a) Profile of the two synthetic surfaces.



(b) Surface of Ganymede, Jupiter's moon.

Figure 2.2: (a) Surface profiles; (b) Surface of Ganymede

<i>Roughness</i>	<i>Sphere</i>	<i>Fractal</i>
$\delta$	22.11	6.39
$R_a$	19.15	4.54
$p_{rms}$	1.34	0.55
$q_{rms}$	1.34	0.55

Table 2.1: Roughness measures for the smooth sphere and the rough fractal

Shape from Shading techniques estimate the surface facet from the image pixel. The slope of a facet is the average slope of the surface area represented by an image pixel. To define a rough surface, we should consider the image resolution. This could be avoided by using rms slope as the roughness measurement. In this thesis, we are restricted to textures whose rms slope is up to 0.5.

## 2.2 Description of textures

Most shape recovery techniques are designed for surfaces like the sphere of the previous example, i.e. smooth objects. We limit our work to the estimation of globally flat, rough surfaces with low slope angles. Furthermore, the surface  $s$  is a potential function to the gradient field  $\mathbf{S}$  which we assume is conservative. It is necessary that the surface can be reconstructed from its gradient field (slopes), i.e. the gradient field must be integrable.

This implies that the components of  $\mathbf{S}$  are continuous and have continuous partial derivatives and that the mixed second derivatives of  $s$  are equal:

$$\frac{\partial^2 s(x, y)}{\partial x \partial y} = \frac{\partial^2 s(x, y)}{\partial y \partial x} \quad (2.6)$$

This is analogous to a smoothness constraint since a surface with discontinuities violates equation 2.6. For the surface to be differentiable into a conservative gradient field, it has to be continuous [22]. Therefore, we restrict our test textures to continuous surfaces.

Phase rich textures show a structured phase spectrum and are characterised by step changes in height, i.e. discontinuities. By assuming surfaces with random phase spectra and band-limited power spectra, we ensure continuity. A further characteristic of random phase test textures is that their height distribution can be assumed to be Gaussian. This is a consequence of the central limit theorem and their broadband spectra. Each frequency component has random phase and its amplitude is therefore a random variable. The sum of many random variables will



have a distribution that converges to the Gaussian. Since differentiation is a linear operation, Gaussian height models give Gaussian slope distributions. The Gaussian assumption is valid for many natural surfaces which are the result of a large number of random effects [23]. Surfaces produced by engineering methods are less likely to possess Gaussian statistics than those arising from natural processes.

Our real test textures are also band-limited. The surface is sampled when it is imaged, the sampling frequency depending on the CCD array and the camera optics. Rapid changes in height might not be seen in the surface image. Thus we assume a surface which is band-limited by Nyquist frequency.

Next we include a description of the synthetic surface models and the real textures that are used throughout the thesis.

## 2.3 Synthetic textures

We use three models of surface topography that are defined in terms of power spectrum. The Power Spectrum Density function (PSD) does not completely define a surface: for that we would also need to know its phase spectrum. From our assumptions, the phase spectra of our models are realisations of an uncorrelated random process.

We chose these surface models because they are well known and allow us to easily modify the surface topography. The software for generating these surface models was written by Dr. G. McGunnigle and Dr. L. Linnet.

Fractals are very frequently used and referred within the vision literature [24] [25] [26] [27]. A fractal Brownian process has shown to accurately model a wide range of natural surfaces [28] [29]. These processes are characterised by their statistical self-similarity, i.e. they are self-affine. A self similar process has the property that a piece of the process,  $1/n^{th}$  the size of the original, when scaled by the factor  $n$ , is in all statistical respects the same as the whole process. For a fractal Brownian process, the scaling factor is not  $n$  but  $n^H$ , with  $0 < H < 1$ . The fractal dimension

is given by  $D = E + 1 - H$ , where  $E$  is the Euclidean dimension of the function. For a two-dimensional process, like a fractal surface, the fractal dimension is constrained to lie somewhere between 2 and 3. The higher the fractal dimension, the rougher the surface. The fractals used in this thesis will have  $D = 2.15$ , since a fractal dimension of 2.15 seems to accurately model many natural surfaces [28]. Importantly, a fractal Brownian model is Gaussian: it is fully specified by its power spectrum. We can safely assume a random phase spectrum since non-random fractals are rarely found in nature [30]. Fractals present a PSD exhibiting an inverse power law with increasing frequency. Assuming that the fractal is spatially isotropic, then its two-dimensional power spectrum is given as a function of the radial spatial frequency, or wavenumber,  $\omega$  by

$$S_f(\omega, \theta) = \frac{k_f}{\omega^\beta} \quad (2.7)$$

Where:

- $S(\omega, \theta)$  is the power spectrum of the surface.
- $k_f$  is a constant.
- $\beta$  is the power roll-off factor and  $\beta = 8 - 2D$ .

Since  $D$  was chosen to be equal to 2.15, the power roll-off factor is 3.7. The roughness of the surface texture will be easily modified by varying the constant  $k_f$ . A rendered fractal with surface roughness equal to 0.25 pixel-width (pw) is shown in Figure 2.3 (left). It is interesting to note how closely it resembles a natural surface.

A more empirical approach to spectral description is to split the spectrum into two different fractal dimensions. Mulvaney *et al.* [31] defined a model for surface profiles that had a flat spectrum at low frequencies, and a roll-off factor of 3.0 for high frequencies. The PSD for the two-dimensional realisation of this type of surface is defined in equation 2.8:

$$S_m(\omega\theta) = k_m \left( \frac{\omega^2}{\omega_c^2} + 1 \right)^{-3/2} \quad (2.8)$$

$\omega_c$  is the cut-off frequency where the inverse power starts to occur and  $k_m$  is a constant related to surface roughness. We set the cut-off frequency to 32 cycles per image (cpi). Figure 2.3 (centre) is a rendered Mulvaney surface with roughness of 0.25 pw. This isotropic surface resembles a rough texture which has undergone some degree of physical processing.

Our third model is described by Ogilvy in her book [23]. She described a directional surface modelled with different power spectrum characteristics for different directions. Directionality is expressed in the cut-off frequency at which fractal behaviour begins to occur for the vertical and horizontal directions. The PSD for this texture is expressed by:


$$S_o(u, v) = \frac{k_o \delta_s^2}{\lambda_1 \lambda_2} \frac{1}{\left(\frac{1}{\lambda_1^2} + u^2\right)} \frac{1}{\left(\frac{1}{\lambda_2^2} + v^2\right)} \quad (2.9)$$

Where:

- $S(u, v)$  is the power spectrum of the directional surface.
- $k_o$  is a constant.
- $\lambda_1^{-1}, \lambda_2^{-1}$  are the cut-off frequencies for the horizontal and vertical directions.
- $\delta_s^2$  is the surface variance.
- $u, v$  are the Cartesian frequencies.

The surface height distribution is Gaussian, with zero mean and standard deviation of  $\sqrt{2}/\lambda_1$  and  $\sqrt{2}/\lambda_2$  for the  $x$  and  $y$  directions. The surface roughness as well as the directionality can be controlled in this model. Figure 2.3 (right) shows an example of an Ogilvy surface with cut-off frequencies of  $\lambda_1^{-1}=32$  cpi and  $\lambda_2^{-1}=16$  cpi, and surface roughness of 0.25 pw. Since it is an anisotropic surface, the rms slope varies with the direction,  $p_{rms}$  and  $q_{rms}$  are 0.13 and 0.10 respectively.

It is necessary here to clarify that either an increase in rms height or a decrease in surface correlation length leads to an increase in rms gradient. Therefore, the



The surfaces have the same reflectance characteristics and they are lit under the same conditions

Figure 2.3: Synthetic textures: Fractal (left), Mulvaney (centre) and Ogilvy (right)

rms gradient or slope is controlled either by the constants  $k_f$ ,  $k_m$  and  $k_o$  or by the roll-off factor and cut-off frequencies.

Ogilvy's model is incapable of generating isotropic surfaces: if both cut-off frequencies are equal, a bidirectional surface is produced. The model does not allow a transition between highly directional and isotropic surfaces. Thus we restrict its use to unidirectional surfaces. We should estimate and avoid the range of cut-off frequencies that produces bidirectional textures.

For a unidirectional surface the polar plot, the distribution of power over direction, shows one single peak. In contrast, the polar plot of a bidirectional surface exhibits two peaks. To assume an unidirectional surface, the lower peak should be at least 3dB smaller than the peak due to the main directionality. Figure 2.4 shows the polar plots of two Ogilvy surfaces. The cut-off frequencies of the first surface are set to the same value (16 cpi) and its polar plot shows two peaks, around  $0^\circ$  and  $90^\circ$ . For the second surface the cut-off frequencies are far apart being  $\lambda_1^{-1} = 512$  cpi and  $\lambda_2^{-1} = 16$  cpi. Its polar plot exhibits one single peak around  $0^\circ$  and this surface can be considered unidirectional.

Table 2.2 shows some examples of Ogilvy surfaces which can be considered unidirectional. The value of  $d$  to assume unidirectionality depends on the surface's cut-off frequencies. For higher frequencies, the directionality condition is relaxed. This investigation shows a deficiency of the model that has to be considered when generating Ogilvy textures. By setting the cut-off frequencies to  $\lambda_1^{-1}=32$  cpi and  $\lambda_2^{-1}=16$  cpi, the directionality  $d$  is equal to 0.57. This assures unidirectionality.




Figure 2.4: Polar plot of two Ogilvy surfaces

$\lambda_1^{-1}$	$\lambda_2^{-1}$	$p_{rms}$	$q_{rms}$	$d$
512	110	0.085	0.082	0.511
256	90	0.091	0.084	0.519
128	58	0.102	0.089	0.534
64	32	0.115	0.094	0.549
32	18	0.120	0.097	0.553
16	9	0.127	0.099	0.560

Table 2.2: Description of 6 Ogilvy textures that comply with the unidirectionality condition

The surface models are band-limited above the Nyquist frequency and below the lowest frequency. For the fractal model to have a finite variance, the surface is band-limited below the lowest frequency. This is because the fractal model is nonstationary. The integral of its PSD is infinite and, therefore, its variance is also infinite. By low pass limiting the surface spectrum, we are making the test sample small compared to the lowest frequency undulation of the surface. Then the nonstationary autocorrelation function (ACF) is almost stationary and can be approximated. The PSD can be derived: it will be stationary except for its apparent steady value which is proportional to the lowest frequency undulation of the surface [32].

## 2.4 Real textures

### 2.4.1 Description

Our real textures have a variety of topographies and reflectance functions. In this section we give a description of the surface's textural characteristics, the processes that formed them and the materials of which they are composed. Our data includes several fractures, deposit surfaces, sand-ripples, anaglyptas, sand papers, textiles. They are imaged under the same lighting conditions in Figure 2.5 and Figure 2.6.

Our database contains five surfaces formed by fracturing large blocks of plaster. These textures differ in their roughness but all have a matte appearance. The *fractures* (a), (b) and (c) are examples of a rough surface, a moderate roughness surface and a gentle roughness surface. For the *fractures* (d) and (e) the plaster blocks were cured resulting in two smooth ceramic type surfaces. Therefore their appearance is slightly more glossy.

The second group of real textures consists of six *deposit* surfaces. In Figure 2.5 (f) and (g), two deposit fractures are imaged. Plaster powder was deposited on a plaster mould. The moulds correspond to a moderate roughness fracture and a rough fracture. For the surfaces (h), (i), (j) and (k), a flat plaster was used as




Figure 2.5: Real textures: (a) rough fracture; (b) moderate roughness fracture; (c) gentle roughness fracture; (d) smooth cured fracture; (e) smooth cured fracture with fracture patterns; (f) moderate roughness deposit fracture; (g) rough deposit fracture; (h) heavy deposited surface; (i) medium deposited surface; (j) light deposited surface; (k) sparse deposited surface and (l) textile.

a base. These textures differ on the amount of deposit, ranging from heavily to sparsely deposited surfaces. We should bear in mind that while the heavy deposits are near Gaussian, the more sparsely deposited surfaces are phase-rich [33]. The *deposit* surfaces are made of plaster, thus their appearance is matte.

Three *sand-ripple* surfaces are also part of our collection of real textures. They were formed by wave action on a plaster surface. The first of these textures is a low frequency, near Gaussian sand-ripple (Figure 2.6 (a)). In contrast, the texture in Figure 2.6 (b) is phase rich. The last sand-ripple is made using a different brand of plaster which appears to contain some quartz-like material (Figure 2.6 (c)). This resulted in a less directional and less matte surface.

The last set of textures consists of four types of *anaglypta* wallpaper. These wallpapers have a plastic-like appearance, thus they are not matte. They resemble an anaglypta formed by *irregular stripes* Figure 2.6 (d), a *net* (e), some *ripples* (f) and some *rice grain shapes* (g). The first three textures are directional whilst the last is isotropic.

The textures in this database are relatively smooth and approximately matte. We should consider surfaces that are either extremely rough or glossy. Also non-constant albedo surfaces or heavily shadowed images should be considered. This will provide our database with textures which would be out of the scope of our work, and allow us find the limits of texture shape recovery. Figure 2.6 shows some of these exceptional textures. The surface imaged in Figure 2.6 (h) is a *sand paper*. It has interesting optical properties, since some facets are highly glossy. As we change the light azimuth, we observe that individual facets reflect intensity over a narrow range of angles - corresponding to a sharp intensity lobe in the reflectance function. Figure 2.6 (i) shows a *repetitive primitive* texture formed by pearl barley. It is a phase rich surface (discontinuous), very rough and presents non-constant albedo. Furthermore, its images are heavily shadowed. Figure 2.5 (l) shows a non-constant albedo *textile*. This is a rough texture with approximately matte reflectance. The *irregular stripes anaglypta* of Figure 2.6 (d) is a very rough surface and its images are






Figure 2.6: Real textures: (a) low frequency sand ripple; (b) high frequency sand ripple; (c) less directional sand ripple; (d) irregular stripes anaglypta; (e) net anaglypta; (f) ripples anaglypta; (g) rice grain shape anaglypta; (h) sand paper; (i) pearl barley

heavily shadowed. Therefore, this anaglypta is also considered among the special textures.

### 2.4.2 Do the real textures comply with the surface assumptions?

We assume that:

1. The textures are globally flat.
2. They are rough.
3. They have low slope angles.
4. They have conservative gradient fields.

We estimate some of the characteristics of our textures to assess the validity of these four assumptions.

A **globally flat** surface has a gradient field with zero mean. We use a classic Photometric Stereo technique to estimate the first order surface derivatives from the surface images (see Section 4.5.1). The mean value of the surface derivatives is near zero for most textures (Table 2.3). We did not remove the mean value of the intensity images before estimating the surface gradient. Removing the intensity mean is approximately like removing the gradient mean without affecting the second and higher order statistics of the surface (see linear approximation of Lambert's law equation 3.12). Thus we could somehow force the surface to be globally flat. For the *smooth cured fracture*, the *low frequency sand ripple* and the *smooth sand ripple* the value of  $\bar{q}$  is 0.103, -0.175 and 0.255 respectively. By removing the mean of the images,  $\bar{q}$  becomes 3.50e-3, 4.27e-4 and -1.68e-4.

We estimate the **rms roughness** ( $\delta$ ) and the **average roughness** ( $R_a$ ) in pixels-width for the real textures (Table 2.3). The gradient fields are integrated into the surface using a frequency integration algorithm (see Section 4.5.2). The fracture surfaces are very rough textures. The deposit surfaces decrease in roughness as the amount of deposit decreases. The sand ripple textures are not rough surfaces. The

anaglyptas show larger values of roughness than the sand ripples. We should note that Table 2.3 shows the value of roughness for the estimates of the real surfaces. The method used for estimating the surface has some limitations (see Section 5.4). The last four textures, the special surfaces, are not expected to be accurately estimated, thus the calculated roughness is not reliable.

To validate the **low slope angles** assumption, we compute the *rms slope* of the textures (Table 2.3). The largest measured rms slope is of 0.406, this translates into an average slope angle of  $22^\circ$ . We can safely assume that the textures present low slopes.

We use two criteria to test whether our real data is **integrable**. To ensure integrability, the mixed second derivatives have to be equal and the surface gradient must be continuous. A Gaussian gradient distribution corresponds to a surface whose facets have low slopes. Therefore the surface facets are likely to be seen from the viewer and the surface gradient is continuous. We assess how Gaussian the surface gradient distribution is and how similar the mixed second derivatives are.

We compare the distribution of the surface gradient with its best fit Gaussian. This fitted Gaussian has the same mean and standard deviation as the slope distributions. It is not the optimal Gaussian fit for the data in the least squares sense, but it allows us to infer how Gaussian the gradient is.

Table 2.4 shows the normalised root mean square error between the gradient and the Gaussian fit computed as:

$$e_{rms} = \frac{\sqrt{\sum_i (f_i - g_i)^2}}{\sqrt{\sum_i (f_i - \bar{f})^2}} \quad (2.10)$$

where  $f$  is the surface gradient,  $g$  is the best Gaussian fit to  $f$  and  $i$  is a point of the surface. The first column in Table 2.4 is the rms error for the  $p$  derivative and the second column for the  $q$  derivative. The third column shows the rms error between the mixed second derivatives, i.e.  $f$  is  $\left(\frac{\partial^2 s}{\partial x \partial y}\right)$  and  $g$  is  $\left(\frac{\partial^2 s}{\partial y \partial x}\right)$  in equation 2.10.

<i>Texture</i>	$\bar{p}$	$\bar{q}$	$\delta$	$R_a$	$p_{rms}$	$q_{rms}$
<i>Rough fracture</i>	0.015	-0.072	14.268	11.865	0.323	0.328
<i>Moderate roughness fracture</i>	-0.010	0.118	7.500	6.413	0.290	0.262
<i>Gentle roughness fracture</i>	-0.073	0.001	9.577	7.449	0.229	0.275
<i>Smooth cured fracture</i>	-0.054	0.103	7.920	6.772	0.158	0.241
<i>Smooth patterned fracture</i>	-0.086	0.012	7.199	5.700	0.242	0.217
<i>Moder rough deposit fracture</i>	0.006	0.060	3.236	2.531	0.251	0.211
<i>Rough deposited fracture</i>	-0.029	0.042	6.675	5.376	0.297	0.282
<i>Heavy deposited surface</i>	-0.092	-0.026	4.115	3.533	0.293	0.305
<i>Medium deposited surface</i>	-0.024	0.001	2.386	1.883	0.256	0.288
<i>Light deposited surface</i>	-0.059	-0.008	0.794	0.633	0.218	0.238
<i>Sparse deposited surface</i>	0.034	-0.040	1.194	0.982	0.231	0.253
<i>Low frequency sand-ripple</i>	-0.022	-0.175	0.975	0.791	0.311	0.163
<i>High frequency sand-ripple</i>	-0.070	0.017	0.559	0.455	0.277	0.080
<i>Smooth sand-rippled surface</i>	-0.074	0.255	1.159	0.944	0.252	0.156
<i>Net anaglypta</i>	0.047	0.039	2.142	1.464	0.370	0.292
<i>Rippled anaglypta</i>	0.021	0.056	1.999	1.604	0.329	0.225
<i>Rice grain shaped anaglypta</i>	0.069	-0.049	1.295	1.017	0.210	0.189
<i>Irregular stripes anaglypta</i>	-0.008	-0.009	2.425	1.916	0.406	0.185
<i>Sand paper</i>	-0.017	0.011	1.211	1.014	0.070	0.068
<i>Repetitive primitive surface</i>	-0.001	0.012	3.900	3.053	0.255	0.302
<i>Textile</i>	-0.044	-0.071	1.373	1.106	0.252	0.237

Table 2.3: Gradient mean and roughness for the real surfaces

<i>Texture</i>	$e_{rms}(\hat{p})$	$e_{rms}(\hat{q})$	$e_{rms}(\hat{p}_y = \hat{q}_x)$
<i>Rough fracture</i>	6.26%	17.34%	32.86%
<i>Moderate roughness fracture</i>	9.52%	5.55%	21.02%
<i>Gentle roughness fracture</i>	9.39%	4.89%	23.51%
<i>Smooth cured fracture</i>	5.28%	9.28%	20.56%
<i>Smooth patterned fracture</i>	4.91%	26.53%	19.98%
<i>Moder rough deposited fracture</i>	5.56%	3.85%	22.54%
<i>Rough deposited fracture</i>	6.62%	9.02%	24.68%
<i>Heavy deposited surface</i>	10.71%	10.32%	22.68%
<i>Medium deposited surface</i>	8.29%	7.71%	23.73%
<i>Light deposited surface</i>	6.01%	5.30%	20.91%
<i>Sparse deposited surface</i>	4.31%	4.52%	23.22%
<i>Low frequency sand-ripple</i>	37.50%	17.61%	42.39%
<i>High frequency sand-ripple</i>	30.76%	51.16%	34.38%
<i>Smooth sand-rippled surface</i>	13.35%	17.57%	31.97%
<i>Net anaglypta</i>	14.53%	35.44%	22.56%
<i>Rippled anaglypta</i>	16.67%	25.89%	21.32%
<i>Rice grain shaped anaglypta</i>	19.81%	24.67%	19.25%
<i>Irregular stripes anaglypta</i>	37.77%	66.27%	34.45%
<i>Sand paper</i>	16.45%	16.57%	65.22%
<i>Repetitive primitive surface</i>	16.40%	13.81%	50.14%
<i>Textile</i>	4.10%	5.21%	58.42%

Table 2.4: Validity of the integrability assumption for the real textures

The *fractures* and the *deposit* surfaces match a Gaussian distribution quite closely giving a rms error smaller than 10% in most cases. Although these distributions are not perfectly Gaussian they are bell-shaped and can be approximated. The *textile* surface also has a gradient distribution which is roughly Gaussian.

In contrast, the *sand ripples*, the *anaglypta* surfaces, the *sand paper* and the *repetitive primitive* texture don't comply with the Gaussian assumption. The rms error between the gradient distribution and its best fitted Gaussian is larger than 10%. The surface gradient is likely to be discontinuous, thus the surfaces violate the integrability condition.



The difference between the mixed second partial derivatives is large for all textures. This is partly due to the non-ideal central difference derivative operator used to estimate the second derivatives (see Section 5.3.2). The error is larger for rougher surfaces, such as the *rough fracture* and for phase rich surfaces, such as the *sand ripples*. We cannot assume that the so-called special textures are integrable, i.e. the *irregular stripes anaglypta*, the *sand paper*, the *repetitive primitive* and the *textile*.

The gradient distributions of some of the real textures are plotted together with their Gaussian fit in Figure 2.7. The selected examples are two Gaussian surfaces: a *fracture* (Figure 2.5 c) and a *deposit* surface (Figure 2.5 j); and two non-Gaussian surfaces: a *sand ripple* (Figure 2.6 a) and an *anaglypta* (Figure 2.6 e). Appendix A contains the slope distributions for all the real textures.



### 2.4.3 Texture sampling

Table 2.5 details the image sampling of our real textures. The *fractures*, the *deposit* surfaces, the *sand ripples* and the *repeating texture* are lit under light zeniths of 45°, 60° and 75°. For each zenith angle, 12 images were captured, i.e. the textures were captured for light azimuth increments of 30°. This gives a total of 36 images per texture.



The *anaglypta* surfaces and the *sand paper* were captured for two zenith angles: 60° and 75°. Similarly, the light azimuth was moved in 30° increments, giving a





(a) Gentle roughness fracture



(b) Light deposited surface



(c) Low frequency sand ripple



(d) Net anaglypta

Figure 2.7: Slope distributions of real surfaces (real and fitted Gaussian)

	$\sigma = 45^\circ$	$\sigma = 60^\circ$	$\sigma = 75^\circ$
<i>Fractures</i>	$\Delta\tau = 30^\circ$	$\Delta\tau = 30^\circ$	$\Delta\tau = 30^\circ$
<i>Deposits</i>	$\Delta\tau = 30^\circ$	$\Delta\tau = 30^\circ$	$\Delta\tau = 30^\circ$
<i>Sand ripples</i>	$\Delta\tau = 30^\circ$	$\Delta\tau = 30^\circ$	$\Delta\tau = 30^\circ$
<i>Anaglyptas</i>	-	$\Delta\tau = 30^\circ$	$\Delta\tau = 30^\circ$
<i>Sand paper</i>	-	$\Delta\tau = 30^\circ$	$\Delta\tau = 30^\circ$
<i>Repetitive primitive</i>	$\Delta\tau = 30^\circ$	$\Delta\tau = 30^\circ$	$\Delta\tau = 30^\circ$
<i>Textile</i>	$\Delta\tau = 30^\circ$	$\Delta\tau = 30^\circ$	-

Table 2.5: Photometric sampling for real textures

total of 24 images per texture.

Finally, the *textile* was lit from zeniths of  $45^\circ$  and  $60^\circ$ ; azimuth was varied in  $30^\circ$  increments. Therefore, 24 images were captured for the *textile*.

## 2.5 Discussion

In this chapter we introduced the synthetic and real textured surfaces that will be used in subsequent chapters. Some definitions relevant to the thesis were also included. The type of surfaces that are within the scope of this work are globally flat rough textures with low slope angles. They should be band-limited continuous functions with conservative gradient fields.

We introduced three surface models: Fractal model [30], Mulvaney’s model [31] and Ogilvy’s model [23]. The models were uniquely described in terms of their power spectrum, since a random phase spectrum was assumed. They are parametric models that allow us to control the surface characteristics. Furthermore they are well known in the literature and extensively used. They are band-limited above the Nyquist frequency (pixelation) and below the lowest frequency. A limitation of Ogilvy’s model was pointed out, and the conditions for using it were reported.

We included a description of real data which covered their topographies, re-



flectance characteristics and degrees of directionality and roughness. Our data base consists of 21 textures that we divide in five groups: *sand ripples*, *fractures*, *deposits*, *anaglyptas* and the special surfaces. Several textures have the same topography and reflectance, but they differ in surface roughness. Others differ on the surface materials, thus they show different reflectance characteristics.

The validity of the surface assumptions was assessed for our real textures. The assessment was four-fold. First, we test whether our real textures are globally flat concluding that we could force the surface estimates to be globally flat. Secondly, we measured the surface roughness of our textures. Thirdly, we proved that the real surfaces have low slope angles. And finally we assessed the validity of the integrability condition. We concluded that the *fractures* and the *deposits* could be considered Gaussian. Furthermore, their mixed second derivatives matched each other with a rms error of 20 to 25%. In contrast, the *sand ripples* and *anaglyptas* were not Gaussian and their mixed second derivatives were not similar in most cases.

# Chapter 3

## Surface-to-Image Models

---

In this chapter we try to model the transfer function from the surface to the image. To understand the nature of this task, some knowledge of the theory of light reflection is necessary. We start the chapter by defining some relevant terms. Then we review the background theory of light reflection. Six main phenomena arise from light-object interaction: reflection, transmission, absorption, diffraction, refraction and interference. Although all these phenomena have been modelled to some extent, in computer graphics most attention has been paid to reflection. We survey reflection models from the literature. The models reviewed in this chapter are *local*, i.e. only direct, or first, reflection of light from a surface facet is considered. No *global*<sup>1</sup> or indirect illumination (reflection from another surface facet) is taken into account.

We choose two of the reviewed reflection models as the rendering models to use in our simulations. In this way, we render matte surfaces using Lambert's law and we simulate surfaces with a plastic-like appearance using Phong's model. Although, strictly speaking, it is not a reflection model, Kube's approximation to Lambert's law is also described.

In this chapter we use simulation to compare the three reflection models. We first study the validity of Kube's model. Secondly, we analyse the similarities between Lambert's law and Phong's model in terms of surface appearance.

---

<sup>1</sup>*Ray tracing* and *radiosity* are established global models that simulate reflection. Ray tracing approximates specular reflection and light transmission, and excludes other considerations. The radiosity method models the interaction of diffusely reflecting surfaces [34].




Figure 3.1: The geometry of the imaging system

### 3.1 Definitions

Some concepts relevant to this chapter are here introduced. We describe our imaging system. We also define shadowing as well as some radiometric terms.

**Our imaging system** is shown in Figure 3.1. The optical axis is aligned with the  $z$  axis of the coordinate system. The reference plane of the surface is perpendicular to the viewing direction. The origin of the system is at the centre of the surface. The viewer is far away from the surface relative to the size of the surface, so that orthographic projection can be assumed. The light source is a point source and is far away from the surface, so that constant illumination over the scene can be assumed. The azimuth ( $\tau$ ) and the zenith ( $\sigma$ ) of the light source are also called illuminant tilt and slant respectively.

**Shadowing** When speaking about reflection, terms such as masking, cast and self shadowing often appear.

*Masking* is used to describe the effect of obstructions in the reflected light (Figure 3.2 a) and *shadowing* to describe obstruction in the incident light (Figure 3.2 b). We can neglect masking effects since the viewer in our imaging system is straight above a globally flat surface.




Figure 3.2: Masking (a) and self and cast shadows (b)

*Self shadows* occur when the difference between the unit vector normal to the surface facet  $\mathbf{N}$  and the illuminant unit vector  $\mathbf{L}$  is larger than  $90^\circ$ . In Figure 3.2 (b), there are two facets that form a V-shaped groove, the facet of the right is self shadowed.

A facet is *cast shadowed* when another facet projects a shadow over it. In Figure 3.2 (b), the lower half of the facet on the left is cast shadowed by the facet on the right.

**The radiometric terminology** found in many text books [18] should be introduced before investigating light reflection. *Irradiance* is the incident light energy per unit surface area ( $Wm^{-2}$ ). *Radiance* is the light energy radiated per solid angle in a particular direction ( $Wm^{-2}St^{-1}$ ). The solid angle  $\Omega$  subtended by a small planar patch of area  $A$  at distance  $r$  is defined in Figure 3.3 (a). *Brightness* is determined by the amount of energy an imaging system receives per unit foreshortened area ( $Wm^{-2}St^{-1}$ ). The foreshortened area ( $A_e$ ) is the surface facet area ( $A$ ) times the cosine of the angle ( $v$ ) between the surface normal and the radiated light, i.e. the viewing direction (Figure 3.3 b).

To estimate the scene radiance, we measure the brightness received by the CCD array. Horn has shown that the image irradiance is proportional to the scene radiance [18].




Figure 3.3: (a) Solid angle; (b) Foreshortened area




Figure 3.4: A BRDF relates incoming energy  $(\phi_i, \theta_i)$  to outgoing intensity  $(\phi_v, \theta_v)$

## 3.2 Theory of light reflection

We describe the interaction of light with a surface by relating the incoming and outgoing radiances at a given point  $P$  on the surface. This expression is known as the Bidirectional Reflectance Distribution Function (BRDF) and is defined as:

$$R_{bd}(\lambda, \phi_i, \theta_i, \phi_v, \theta_v) \quad (3.1)$$

where  $\lambda$  is the wavelength of the incident light,  $(\phi_i, \theta_i)$  is the direction of the incoming light and  $(\phi_v, \theta_v)$  is the direction of the outgoing light (Figure 3.4).

In general, light may be reflected in an orientation or wavelength dependent way; thus, the BRDF may be very complex. Koenderink and van Doorn [35] have measured the BRDFs for a variety of surfaces. Their measurements indicate that

BRDFs are fairly well behaved and can be represented using a relatively small (50 or less) number of basis BRDFs. Even 50 parameters makes the problem of surface estimation from images intolerably hard.

Dana obtained parametric descriptions of the BRDFs compiled on her database [36]. This consisted of reflectance measurements for over 60 very different samples, each observed with over 200 different combinations of viewing and source directions. Two attempts were made to fit Koenderink's representation [37] to the BRDF measurements. This resulted in a 5 parameter and a 55 parameter representation, respectively. She concluded that the simpler representation was insufficient for modelling all samples in the database, unlike the 55 parameter representation.

Recent work attempts to reconstruct the geometry of objects with arbitrary and possibly anisotropic BRDFs [38]. The proposed methods make no assumption about the object's shape, the presence of shadows, or the nature of the BRDF which may vary over the surface. But again, a large collection of images from the object is needed (up to 143 images), making this method time and resource consuming, beyond the scope of this thesis.

Using the BRDF leads to overly complicated models with too many parameters. However, accurate models of surface BRDFs are simplified into reflectance models used in computer vision. There are two simplifications we can make to the BRDF. The first is almost universal in computer graphics and vision: assuming that the microstructure of the surface facet is isotropic, the two azimuth variables may be replaced by their difference  $\phi_i - \phi_v$ . The second simplification is due to our use of a monochrome camera - this integrates irradiance over wavelength and we drop the wavelength variable.

Reflectance models only consider the reflection and transmission of light. The degree to which either occurs depends on the electrical characteristics of the material. In this way, reflection is dominant in conductors and transmission in dielectrics, although both exhibit transitional behaviour. The reflection associated with conductors is more directional than the reflection from dielectrics which is due to a




Figure 3.5: Reflection: (a) mirror-like specular (b) specular+diffuse (c) perfectly diffuse

different mechanism known as body scattering. For a flat conductor, the light is reflected in the mirror direction:

$$\phi_v = \phi_i \quad \theta_v = \theta_i \quad (3.2)$$

This is known as the specular contribution. The smoother the surface the larger this contribution. For a perfect mirror, the specular reflection would be the sole contribution (Figure 3.5 (a)). For a rougher surface, the specular contribution would be spread over a larger range of angles (Figure 3.5 (b)). This is translated into a more diffuse reflection which exhibits a specular peak. The ideal diffuse contribution reflects uniformly in all directions into the hemisphere (Figure 3.5 (c)). It is only present in dielectric surfaces which we term *Lambertian surfaces*.

### 3.3 Lambert's law

The simplest model, formulated by Lambert in 1760 [20], is the quantitative law for perfectly diffuse surfaces. Lambert stated that a perfectly diffuse surface appears equally bright from all viewing directions.

He assumed that the diffuse surface has a homogeneous reflectance function. He ignored self or cast shadows, as well as any inter-reflection. Under these assumptions and for a point light source at infinity, the reflected intensity from one point on the surface is proportional to the angle between the surface normal at that point and the illuminant direction:

$$\mathbf{I} = I_0 \rho \cos \gamma = I_0 \rho (\mathbf{N} \cdot \mathbf{L}) \quad (3.3)$$

or in scalar form:

$$i(x, y) = i_0 \rho(x, y) \frac{-p(x, y) \cos \tau \sin \sigma - q(x, y) \sin \tau \sin \sigma + \cos \sigma}{\sqrt{p^2(x, y) + q^2(x, y) + 1}} \quad (3.4)$$

where:

- $\mathbf{I}$  or  $i(x, y)$  is the intensity reflected at the point  $(x, y)$ .
- $I_0$  or  $i_0$  is the incident intensity.
- $\rho(x, y)$  is the albedo; a coefficient that represents the proportion of light reflected from the point  $(x, y)$  with respect to the incident light.
- $\mathbf{N} = \left( \frac{-p}{\sqrt{p^2+q^2+1}}, \frac{-q}{\sqrt{p^2+q^2+1}}, \frac{1}{\sqrt{p^2+q^2+1}} \right)$  is the unit vector normal to the surface at  $(x, y)$ .
- $\mathbf{L} = (\cos \tau \sin \sigma, \sin \tau \sin \sigma, \cos \sigma)$  is the unit vector pointing at the light source.
- $\tau, \sigma$  are the light azimuth and zenith defined as in Figure 3.1.
- $\gamma$  is the angle between the normal to the surface at  $(x, y)$  and illuminant direction.

The partial derivatives of the surface  $s$  at the point  $(x, y)$  are defined as:

$$p(x, y) = \frac{\partial s(x, y)}{\partial x} \quad (3.5)$$

$$q(x, y) = \frac{\partial s(x, y)}{\partial y} \quad (3.6)$$

Wolff has demonstrated that the Lambertian model only really applies when the angle of incidence and the angle of reflection are small relative to the surface normal [39]. Our textures have low slope angles and the viewer is directly above the surface, thus the angle of reflection is small. In contrast, the angle of incidence is often large and the Lambertian model would become inaccurate. Wolff has developed a



simple modification of Lambert’s law that accurately account for all illumination and viewing directions.

Oren and Nayar propose a generalisation of Lambert’s model to include non-perfectly diffuse rough surfaces [40]. In their work the surface reflection also depends on both the incident direction and the viewing direction. This model accounts for more complex geometric and radiometric phenomena such as masking, shadowing and inter-reflections between points on the surface, although this last contribution is almost generally neglected.

Although Lambert’s law was advanced almost 250 years ago, it remains one of the most widely used models in machine vision. It is used explicitly by shape recovery techniques, such as Shape from Shading and Photometric Stereo (see Chapter 4). It is also invoked by vision techniques such as binocular stereo and motion detection to solve the correspondence problem [41] [42]. In the field of remote sensing, the Lambertian model is often used to apply brightness corrections to images of the same scene obtained under different illumination conditions [43].

In this thesis Lambert’s law will be often used as our reflectance model. It is a simple model, well known in the field and the most frequently employed in computer vision. It does reasonably well in approximating reflection for a wide range of matte surfaces.

### **3.4 Modelling specularities: Phong’s**

After accounting for diffuse reflection, we now introduce a more complete and more realistic model [44]. Phong’s model is empirical, but it has found great acceptance within the computer graphics community and has become the industry standard [34]. The model is a linear combination of three components: a diffuse, a specular and an ambient component. Some assumptions by the model are:

- Only primary reflection of light is considered; it is a local model.
- The object imaged has uniform reflectance properties.

- The diffuse and specular terms are modelled as local components.
- The viewer and the point light source are at infinity.
- Shadows are ignored.
- The image irradiance which is not due to the primary reflection is modelled as a constant, the ambient component.

We don't consider the ambient component in our implementation of Phong's model. The diffuse contribution is the ideal described by Lambert, and it is evaluated as:

$$I_d = I_0 K_d \cos \gamma = I_0 K_d (\mathbf{N} \cdot \mathbf{L}) \quad (3.7)$$

where  $K_d$  is a wavelength-dependent empirical reflection coefficient, equivalent to the albedo  $\rho$  in Lambert's law.  $K_d$  gives the proportion of light diffusely reflected by the surface facet. The diffuse reflection is a function of  $\gamma$ , the angle between the light unit vector  $\mathbf{L}$  and the unit vector normal to the surface facet  $\mathbf{N}$ . Similarly, the specular component will be a function of  $\eta$ , the angle between the viewing direction  $\mathbf{V}$  and the mirror direction  $\mathbf{R}$  (see Figure 3.5):

$$I_s = I_0 K_s \cos^n \eta = I_0 K_s (\mathbf{R} \cdot \mathbf{V})^n \quad (3.8)$$

Both  $\mathbf{V}$  and  $\mathbf{R}$  are unit vectors.  $K_s$  is an empirical specular reflection coefficient and denotes the amount of light specularly reflected. The index  $n$  simulates surface roughness, it controls the width of the specular peak. For a perfect mirror  $n$  would be infinite and the reflected light would be constrained to the mirror direction. For glossy surfaces, the value of  $n$  is large; to model the reflection from a smooth surface the lobe is narrow around the mirror direction. For more diffuse reflection, the value of  $n$  is small, giving a wide reflection lobe.

The vector  $\mathbf{R}$  is computationally expensive to calculate. Phong proposed a way to estimate it in [44]. Instead we evaluate the specular term using a common approximation:  $\mathbf{H}$  is a normal vector to a hypothetical surface that is oriented in a




Figure 3.6: Mirror direction for Phong's model

direction halfway between the light direction and the viewing direction (Figure 3.6). The specular term now becomes a function of  $(\mathbf{N} \cdot \mathbf{H})$ . The vector  $\mathbf{H}$  denotes the required orientation for a surface to reflect light maximally along the viewing direction and can be estimated as:

$$\mathbf{H} = \frac{(\mathbf{L} + \mathbf{V})}{2} \quad (3.9)$$

Since the light source and the viewer are considered to be at infinity, the two unit vectors  $\mathbf{L}$  and  $\mathbf{V}$  are constant throughout the scene, and so is  $\mathbf{H}$ . The angle ( $\nu$ ) between  $\mathbf{R}$  and  $\mathbf{V}$  is twice the angle ( $\epsilon$ ) between  $\mathbf{N}$  and  $\mathbf{H}$  (Figure 3.6). This can be compensated for by adjusting the value of  $n$ . Phong suggested  $n$  to vary from 1 to 10, but for our implementation  $n$  varies from 1 to 25.

Adding the diffuse and the specular components together gives the model's expression:

$$I = I_0(K_d(\mathbf{N} \cdot \mathbf{L}) + K_s(\mathbf{N} \cdot \mathbf{H})^n) \quad (3.10)$$

where  $K_d + K_s = 1$ .

The reflected intensity is solely a function of the surface normal  $\mathbf{N}$  and the wavelength for a particular light source.  $\mathbf{L}$ ,  $\mathbf{V}$  and  $\mathbf{H}$  are constant.

Phong's model presents some limitations. For instance, the dominant effect of varying  $n$  is not to modify the surface "glossiness", but to make it look as if the size of the light varies. Decreasing  $n$  makes the source look bigger. Thus facet roughness

is not well modelled by the specular term.

Phong’s model is unable to model metallic surfaces; rendered surfaces have a plastic-like appearance.

Another limitation of the model is the neglect of shadows. The absence of cast shadows makes objects appear to be floating above, instead of resting on, a surface. The absence of self shadows makes concavities appear erroneously shadowed.

Since this work is limited to fairly diffuse low slope surfaces, we consider Phong’s model sufficiently good for our rendering purposes. Furthermore, it is popular, simple to implement and the characteristics of the reflected light are easily controlled.

### 3.5 Physically based models

This section reviews reflection models that account for specularities in a more physical way. We survey more recent research in the field and more elaborate models. This survey is included for completeness: it contributes to the general overview on reflectance analysis and synthesis.

A reflection model based on geometrical optics was first introduced by Torrance and Sparrow (T-S) in 1967 [11]. The T-S model is only valid when the surface roughness is large compared to the radiation wavelength. Roughness is modelled by considering the surface as a large number of mirror-like microfacets in the form of symmetric V-shaped grooves. The specular reflection due to the mirror-like facets is a function of the angles of reflection. The diffuse component originates from multiple reflections among the facets or internal scattering. The model accounts for shadowing and masking by adjacent facets.

The T-S model has been widely adopted for modelling specular reflection from rough surfaces. Healey and Binford [45] used the model for computing object surface curvature information. Ikeichi and Sato [46] and Solomon and Ikeuchi [47] used it to estimate surface shape, albedo and roughness. Solomon and Ikeuchi [48] also proposed a non-contact method for measuring surface shape and roughness assuming

a simplified form of T-S model. Wolff and Boult [49] used the model to develop a polarisation stereo method. They presented a polarisation reflectance model based on Fresnel reflection coefficients that accurately predicted the magnitudes of polarisation components of reflected light.

Blinn [50] used a more elaborate scheme than the T-S model. This scheme incorporated a term which accounted for attenuation due to the self-shadowing.

Cook and Torrance's model [51] is based on Blinn's model but assumed a distribution function proposed by [52] that gave the proportionate area of microfacets oriented at a certain angle to the average normal of the surface. This distribution depended on the rms slope of the surface.

Nayar *et al.* [53] proposed a hybrid model which had the simplicity of a geometrical model as well as the completeness of a physical model. They approximated the diffuse lobe by Lambert's law, the specular lobe by T-S model and the specular spike by Beckmann-Spizzichino model [52].

More recently, Lundberg *et al.* [54] criticised T-S model for being in qualitative agreement with specular reflection from rough surfaces but quantitatively inaccurate. They proposed a geometric reflectance model that assumed a microfacet probability distribution which accounted for self-shadowing.

Cook and Torrance's model is similar to Phong's since it is local and divides reflection into three terms. In contrast, it can render polished metallic surfaces correctly. Furthermore, it is accurate for illumination at low angles of incidence unlike Phong's. The main advantage of Cook and Torrance's model is the estimation of colour in the highlight term, but we only consider monochrome images. However, it is more complex and computationally expensive.

## 3.6 Modelling shadows

In this section we explain how we model self and cast shadows. Shadows are modelled and added to the surface image after rendering.

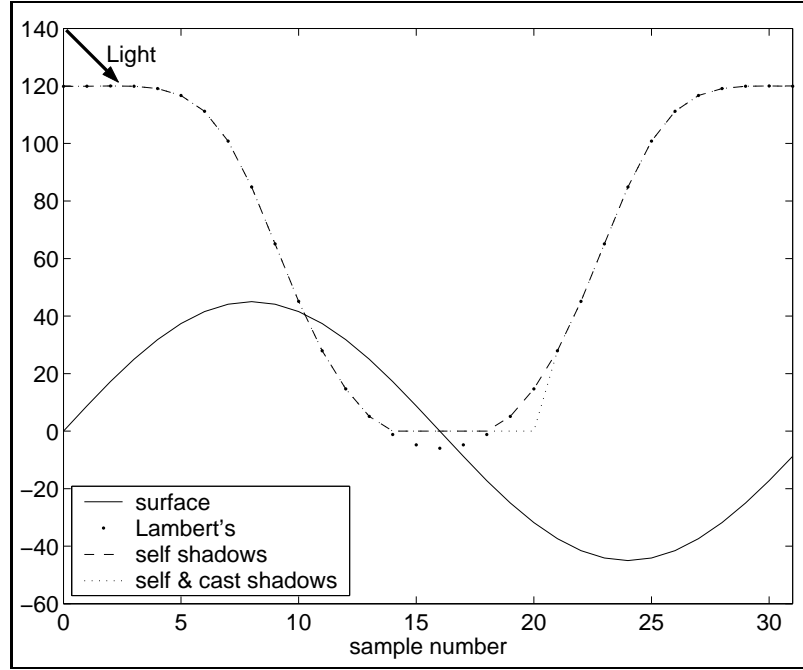


Figure 3.7: Example of self and cast shadows in a sinusoid

The reviewed reflection models predict a negative intensity for a self shadowed facet. By clipping at zero we make all shadows have equal intensity (no light). We make simple geometric calculations to estimate which facets are cast shadowed by others, and their intensity is set to zero. This is unrealistic since although the intensity is small in a shadowed area, it varies due to inter-reflections and ambient illumination. We neglect these two effects, thus shadows are modelled by an absence of light.

An example of shadowing effects is given graphically in Figure 3.7. We render a corrugated sinusoid using Lambert's law. The light has azimuth perpendicular to the directionality of the sinusoid and zenith of  $45^\circ$ . In Figure 3.7 we plot a profile of the surface (continuous line) and its image (unconnected points). The facets from the 14th to the 18th samples are self shadowed and have a negative intensity ( $(\mathbf{N} - \mathbf{L}) \geq 90^\circ$ ). We then plot the same profile after modelling self shadows (dashed line). The intensity of the mentioned samples is been set to zero. Finally, we model cast shadows for the profile (dotted line). The intensity samples 19 and 20 are zero due to the projection of a shadow from some part of the surface over these two facets.

### 3.7 Kube's model

Kube's model is not strictly a reflection model, it is an approximation to Lambert's law designed to allow frequency domain analysis of the rendering process. We present this analytical linear surface-to-image model since it is relevant.

Kube and Pentland [55] linearised the expression in equation 3.4 applying a Taylor's series about  $(p, q) = (0, 0)$ . This leads to the expression shown below:

$$i = i_0 \rho [-p \cos \tau \sin \sigma - q \sin \tau \sin \sigma + \cos \sigma] \left[ 1 - \frac{1}{2}(p^2 + q^2) + \frac{9}{4!}(p^2 + q^2)^2 \dots \right] \quad (3.11)$$

If the surface slope angles are low then  $p \gg p^2$  and  $q \gg q^2$ , i.e.  $p \ll 1$  and  $q \ll 1$ . For low slope angles and  $\sigma > 6^\circ$  a linear approximation to equation 3.11 can be made, the second and higher order terms may be discarded. This leaves only the first three terms:

$$i = i_0 \rho [-p \cos \tau \sin \sigma - q \sin \tau \sin \sigma + \cos \sigma] \quad (3.12)$$

In the frequency domain, the surface derivatives follow the next expressions:

$$P(\omega, \theta) = j\omega \cos \theta S(\omega, \theta) \quad (3.13)$$

$$Q(\omega, \theta) = j\omega \sin \theta S(\omega, \theta) \quad (3.14)$$

where  $j$  represents a phase shift of  $90^\circ$ , and  $\omega, \theta$  are the polar coordinates. Ignoring the mean term and using expressions 3.13 and 3.14, the intensity function of equation 3.12 follows

$$I(\omega, \theta) = -j\omega \cos \theta \cos \tau \sin \sigma S(\omega, \theta) - j\omega \sin \theta \sin \tau \sin \sigma S(\omega, \theta) \quad (3.15)$$

in the frequency domain, which using trigonometry simplifies to:

$$I(\omega, \theta) = -j\omega \sin \sigma \cos(\theta - \tau)S(\omega, \theta) \quad (3.16)$$

This linear function relates the spectrum of the image  $I(\omega, \theta)$  to the spectrum of the surface  $S(\omega, \theta)$ . The image spectrum is a function of both the illuminant azimuth  $\tau$  and zenith  $\sigma$ . Equation 3.16 shows that the imaging process acts as a directional high pass filter. We only need to know the illuminant conditions to recover the surface from its image. Pentland [56] did exactly this using a single image of the surface. He implemented a Wiener filter to attenuate noise and non-linear components of the image intensity pattern.

Chantler [57] investigated Kube's model validity. He divided equation 3.16 in three components, a surface response component, an azimuth response component and a zenith response component. He stated that the surface response component was valid when shadowing occurs and over a range of surface variances. He proposed a raised cosine function of the form  $a + b \cos(\theta - \tau)$  as a more accurate estimate for the azimuth response. Finally, he confirmed that shadowing severely affected the  $\sin \sigma$  relationship.

McGunnigle and Chantler [58] also evaluated Kube's imaging model. The paper assessed the validity of the linear approximation and compared it to the optimal linear filter in terms of the image appearance. They confirmed that the model was valid for surfaces with Lambertian reflection and moderate roughness of *rms slope*  $< 0.3$ . They argued that by introducing a roughness dependent correction factor, the model would be extensive to a large class of rough surfaces. The paper concluded that Kube's model was not optimal, but still valid for relatively rough diffuse surfaces.

### 3.8 Assessment of reflectance models

The assessment is two-fold: first, Lambert's law is compared to Kube's model, and secondly, Lambert's law is compared to Phong's model. With the first simulation, we



investigate the validation of the linear approximation. With the second, we study the contribution of the specular term to the total reflection in terms of surface appearance. For this purpose, Phong’s model is interpreted as a Lambertian model with an added specular term.

In the simulations, we use the synthetic textures described in Chapter 2, i.e. the fractal, the Mulvaney and the Ogilvy surfaces. We adopt a nominal minimal accuracy of 10dB which translates into an error of 10%.

### 3.8.1 Validity of Kube’s model

#### Introduction

Kube and Pentland [55] stated the conditions to discard second and higher order terms in the Lambertian equation. In terms of surface roughness, the condition is  $p_{rms} < 0.1$  and  $q_{rms} < 0.1$ . Since the second order term is proportional to  $p^2 + q^2$  (equation 3.11), the condition becomes  $(p^2 + q^2) < 0.02$ . Third and higher order terms are small relative to the second term, and insignificant relative to the linear term.

A second condition is stated in terms of the light source position: the light zenith has to be larger than  $6^\circ$ . In this thesis the illuminant zenith angle is always larger than  $45^\circ$ . Thus this condition does not need any validation or further investigation.

Chantler [57] studied the validity of Kube’s model by comparing the magnitude of the linear term to the magnitudes of the non-linear terms. He stated that for an error of 3.5%, the rms slope had to be less than  $15^\circ$ . This means that for a signal-to-residue ratio (SRR) of 14.56dB, the rms slope should be less than 0.27. For the threshold of 10dB the restriction of the rms slope can be relaxed to angles larger than  $15^\circ$ .

McGunnigle [59] fitted the reflectance function to its optimal linear approximation. He compared the filter performance for rendering purposes with an empirically derived reflectance map. Unlike Kube’s prediction, his filter was optimal. Neverthe-

less his investigation of the validity of the filter is relevant. He concluded that the SRR was above 10dB, for an isotropic surface with  $p_{rms} < 14^\circ$ , i.e.  $p_{rms} < 0.25$ . He stated that for directional surfaces, the accuracy of the approximation was highly dependent on angle between the illuminant direction and the surface grain direction. He argued that for an accurate rendering the azimuth angle had to be at least  $15^\circ$  apart from the material grain direction.

## Experimental method

We too assess Kube's approximation against surface roughness in the range:

$$(p^2 + q^2) \in [0.02, 0.45, 0.08, 0.125, 0.18, 0.245, 0.32, 0.405, 0.5] \quad (3.17)$$

We render the surfaces with both Lambert's law and Kube's model. The quantity we measure to assess the approximation validity is the signal-to-residue ratio defined in equation 3.18:

$$\frac{S}{R_{lk}} = 10 \log \left( \frac{\text{var}[I_{Lambert}]}{\text{var}[I_{Lambert} - I_{Kube}]} \right) \quad (3.18)$$

where  $I_{Lambert}$  and  $I_{Kube}$  are Lambert's and Kube's images respectively, and  $\text{var}[x]$  is the variance of  $x$ . Four images are rendered with each reflectance model for a fixed zenith angle ( $45^\circ$ ) and azimuth angles in  $90^\circ$  increments. The  $S/R_{lk}$  ratios are then averaged across the four images and the three surfaces to give a more robust assessment.

## Experimental results

The  $S/R_{lk}$  ratio is computed for each pair of Lambertian-Kube's images. The signal-to-residue ratios per surface roughness are averaged and plotted in Figure 3.8. Following Kube's prediction, as the surface roughness increases the approximation diverges from Lambert's law. The  $S/R_{lk}$  ratio is above the 10-dB threshold for  $(p^2 + q^2) < 0.2$ . For an isotropic surface that corresponds to a rms slope of




Figure 3.8: Validity of Kube’s approximation to Lambert’s law.  $S/R_{lk}$  between the intensity images rendered with Lambert’s law and Kube’s model

$p_{rms} = q_{rms} < 0.3$ . This result is in accordance with Chantler’s investigation and McGunnigle’s empirical validation. Kube was more strict with his own approximation, although he did not argue his reasons for the condition on surface roughness. The novelty of our result is that it is been generalised for anisotropic surfaces, the condition is in terms of a non-directional quantity  $(p^2 + q^2)$ .

It is worth noting that the illuminant direction is a key factor for the validity of Kube’s linearisation on anisotropic surfaces. This will be investigated in Chapter 6, Section 6.2.1.

### 3.8.2 What is the effect of specularities?

#### Introduction

Our implementation of Phong’s model consists of two terms, a Lambertian component and a specular component. We here investigate how the two components contribute to the surface appearance. We assess the contribution of the specularity by comparing Phong’s image with a Lambertian image of the same surface.

#### Experimental method

Two experiments are carried out. With the first experiment, we assess the effect of increasing the specular contribution. In the second, the nature of the specularity is




Figure 3.9: Phong's specular term contribution to total reflection against  $K_d$  for  $n = 5$

modified.

The rms roughness of the surfaces is fixed at 0.02 pw. The signal-to-residue ratio is computed for each texture as:

$$\frac{S}{R_{lp}} = 10 \log \frac{\text{var}[I_{Lambert}]}{\text{var}[I_{Lambert} - I_{Phong}]} \quad (3.19)$$

where  $I_{Phong}$  is Phong's image. To improve the assessment, we render four images with each reflectance model for a fixed zenith angle and azimuth angles in  $90^\circ$  increments. We average the  $S/R_{lp}$  ratios across the three textures.

## Experimental results

The strength and shape of the specular peak depend on the values of  $K_d$  and  $n$  respectively (equation 3.10). For the first test, the  $S/R_{lp}$  ratios are plotted against the coefficient  $K_d$  in Figure 3.9. Decreasing  $K_d$  is equivalent to increasing the proportion of specular reflection. The exponent  $n$  is set to 5, i.e. the specular peak is wide but distinguishable.

As  $K_d$  decreases, Phong's image diverges from the Lambertian image, and the  $S/R_{lp}$  ratio decreases. The contribution of the diffuse reflection has to be larger than 75% for a surface to be considered diffuse ( $S/R_{lp} > 10\text{dB}$ ). We only need to render relatively diffuse surfaces, thus we limit  $K_d$  to larger values than 75% for  $n=5$ .




Figure 3.10: Phong's specular term contribution to total reflection against  $n$  for  $K_d = 90\%$

In a second experiment, we measure the  $S/R_{lp}$  ratio for an increasing  $n$ , when  $K_d$  is set to 90% (Figure 3.10). For  $n < 4$  the specular peak is very wide, similar to the diffuse reflection. Phong's image is similar to Lambert's image, and the  $S/R_{lp}$  is high.

As  $n$  increases, the specular peak gets narrower. In the range  $n \in [4, 12]$ , the specular reflection is very distinctive from the diffuse reflection and the  $S/R_{lp}$  ratio decreases almost 10dB.

For large values of  $n$  ( $n > 12$ ) the  $S/R_{lp}$  ratio steadily increases. This is because the specular peak does not affect the surface appearance. The illuminant zenith is set to  $45^\circ$  - in subsequent work  $\sigma \geq 45^\circ$ . The specular reflection follows the mirror direction which is also  $45^\circ$ . Since the peak is very narrow and the viewer is straight above the surface, almost none of the specular reflection reaches the viewer. Thus the surface appearance is similar to Lambert's image, and a higher  $S/R_{lp}$  ratio is measured. We should note that for a near Lambertian surface ( $K_d = 90\%$ ), the  $S/R_{lp}$  ratio is over 10dB for any shape of the peak.

The factor that mainly alters the surface image is the proportion of specularly reflected light. The nature of the specularity is less important for the surface appearance.

## 3.9 Discussion

This chapter gave an introduction to the surface reflectance models that will be used in subsequent chapters. It also included relevant definitions.

We investigated the surface-to-image transfer function. The background theory of light reflection was reviewed. We surveyed the existing reflectance models and described in detail the most appropriate models for our simulation. Three reflectance models, Lambert's law, Kube's approximation [55] and Phong's model [44], were chosen to be thoroughly studied.

We tested the validity of Kube's approximation concluding that the second order term  $p^2 + q^2$  had to be less than 0.2. Importantly, we restricted the light zenith to angles larger than  $6^\circ$ . The contribution of Phong's specular term to the total reflection was also analysed. We concluded that the factor which mainly alters the surface appearance is the proportion of specularly reflected light. It affected the surface image more obviously than the nature of the specularity or the shape of the specular peak.

# Chapter 4

## Shape Recovery: A Review

---

We consider the problem of estimating the shape of rough surfaces from their appearance. This chapter presents a literature survey of existing Shape from Shading techniques. We pay special attention to one class of methods: Photometric Stereo, and identify the advantages of this method for texture estimation. This technique estimates the surface gradient; because our ultimate goal is surface height, we need to integrate the gradient estimate. Several integration techniques are surveyed in this chapter. We conclude that global integration is more appropriate for our purposes.

Finally, we describe a shape estimator that overcomes the difficulties of recovering rough surfaces. It combines a three-light Photometric Stereo algorithm with an efficient global integration method. The combination of these two standard components into a recovery technique for rough surfaces is novel. We describe the implementation of the recovery technique since it is used as a point of comparison in subsequent chapters.

## 4.1 Shape from Shading

### 4.1.1 Introduction

Estimating the shape of a surface from its images is an old problem in computer vision. Horn [18] expresses this problem mathematically through a first-order partial differential equation. Assuming that the surface lies on the  $x - y$  plane, he seeks the surface depth  $s$  in the direction  $z$  that satisfies the image irradiance equation:

$$R(p(x, y), q(x, y)) = E(x, y) \quad (4.1)$$

where  $R(p, q)$  is the reflectance map and  $E(x, y)$  is the image formed by orthographic projection of light onto a plane parallel to the surface plane. The reflectance map expresses irradiance, i.e. the measured intensity, as a function of the facet slope ( $p$  and  $q$ ). Its form depends on the material reflectance characteristics and the illuminant conditions.

This expression for image irradiance relies on two assumptions:

1. The scene radiance depends only on the reflectance of the surface, its normals and the lighting conditions. This implies that the light source is infinitely far away and the inter-reflections between surface facets are ignored.
2. The image irradiance corresponding to a surface point  $(x, y)$  is equal to the scene radiance (see Chapter 3 for definitions).

An interesting case corresponds to the situation where a distant point source illuminates a Lambertian surface. Then according to Lambert's law, the emitted radiance and the reflectance map are given by the cosine of the angle hold between the surface normal direction and the source direction. The irradiance equation for this case takes the form of expression 3.4.



### 4.1.2 Review of single-light Shape from Shading techniques

If the surface is recovered from one single image, the problem of surface estimation is ill-posed. Equation 3.4 has to be solved for three unknowns: the surface derivative fields ( $p$  and  $q$ ) and the albedo. Several single-image Shape from Shading approaches have been proposed. They can be divided into three groups: propagation approaches, minimisation approaches and linear approaches. *Propagation approaches* propagate the shape information from a set of surface points (e.g. singular points) to the whole image. *Minimisation approaches* rely on constraints on the shape of the surface. They obtain a solution by minimising an energy term which is a function of the validity of the constraints for the actual estimate. *Linear approaches* compute the solution based on the linearisation of the reflectance map. In this section we also review the conditions for existence and uniqueness of the solutions.

#### Propagation approaches

Horn's characteristic strip method [18] is essentially a propagation method. A characteristic strip is a line in the image along which the surface depth and orientation can be computed if these quantities are known at the starting point. Horn's method constructs initial curves around the neighbourhood of singular points (points with maximum intensity). Oliensis and Dupuis [60] [61] presented a propagation scheme that formulates Shape from Shading as an optimal control problem. They solve for the surface using numerical methods. Bichsel and Pentland [62] simplified the previous approach and proposed a minimum downhill approach for rapid convergence. Similarly, Kimmel and Bruckstein [63] reconstructed the surface through layers of equal height contours from an initial closed curve.

#### Minimisation approaches

Minimisation approaches produce the optimal solution in terms of the adopted constraints about the surface. Ikeuchi and Horn [64] introduced two constraints: based

on brightness and smoothness. The brightness constraint requires that the reconstructed shape produces the same brightness as the input image. The smoothness constraint ensures an estimated surface which is smooth. The shape of occluding boundaries was given for the initialisation. Brooks and Horn [65] minimised a term which was a function of the same two constraints. They optimised the solution in terms of surface normals instead of height. Frankot and Chellappa [66] enforced integrability in Brooks and Horn’s algorithm in order to recover surfaces with conservative gradient fields. These surfaces have equal mixed second partial derivatives, i.e. they comply with the integrability constraint. More recently Horn [67] replaced the smoothness constraint in his approach with an integrability constraint. The major problem with his method is its slow convergence. Szeliski [68] proposed a faster technique based on a gradient descent algorithm.

Instead of a smoothness constraint, Zheng and Chellappa [69] introduced an intensity gradient constraint. This assures that the intensity gradients of the reconstructed image and the input image are close to each other on both the  $x$  and  $y$  directions. Zhang and Shah [70] also used the intensity gradient constraint. They minimised a function based on the intensity gradient and the smoothness constraints. Leclerc and Bobick [71] solved directly for depth. They used a discrete formulation which relies on smoothness and brightness constraints, and a stereo depth map as a initial estimate. Lee and Kuo [72] also solved for depth using brightness and smoothness constraints without requiring initialisation. They approximated the surface as a union of triangular patches. Wei and Hirzinger [73] introduced a model of the surface depth using radial basis functions that deformed by adjusting their centres, widths and weights to minimise errors in the intensity prediction. They initially used a smoothness constraint to stabilise the solution. A more recent Shape from Shading algorithm satisfied the image irradiance equation as a hard constraint [74]. It used curvature information to impose topographic constraints on the needle map and a robust error function to impose consistency. Cho and Chow [75] proposed an approach that minimises an error intensity function with respect to the weights of

a neural network which represents the 3D object.

### **Linear approaches**

Pentland [56] used a linear approximation to the reflectance function in terms of surface gradient. He applied a Fourier transform to the linear function to get a closed form solution for the depth. Tsai and Shah [76] employed discrete approximations for the surface normal using finite difference of depth and then linearised the reflectance function in depth directly. Their algorithm recovered the depth at each point using an iterative scheme. They subsequently improved their algorithm and extended its use to variable albedo surfaces [77]. More recently Kozera and Klette [78] applied a finite difference-based algorithm to recover the shape of a smooth surface for which the reflectance map was linear. Lee and Kuo’s minimisation approach also assumed a linear reflectance of the surface [72].

### **Existence and uniqueness for single-light Shape from Shading**

To estimate the surface gradient from equation 3.4, we need to solve an unconstrained problem. This raises issues of the existence and uniqueness of the solutions. *Existence* depends on whether a given shading pattern is generated by a true Lambertian surface. *Uniqueness* is whether a shading pattern is due to one and only one Lambertian shape. In the literature, the existence and uniqueness of the solution are mostly investigated for the special case of a surface lit by a single source directly overhead [79], [80], [81], [82] and [83].

One possible solution for the proposed unconstrained nature of the problem is to use more than one surface image. This approach is called Photometric Stereo and is the subject of Section 4.2.

## 4.2 Photometric Stereo

### 4.2.1 Introduction

Woodham [84] states that the shape of a Lambertian surface is uniquely determined by a triplet of images obtained by illuminating a given scene from three different light-source directions. This technique is known as three-light Photometric Stereo, and the solution is determined by a system of the three equations with the form of expression 3.4.

This equation system is solved for three unknowns, the derivative maps of the surface  $p(x, y)$  and  $q(x, y)$ , and the albedo map  $\rho(x, y)$ . 3-light Photometric Stereo techniques don't rely on smoothness constraints. Thus they present a valid approach for shape estimation of rough surfaces and should be further studied.

### 4.2.2 Review of Photometric Stereo techniques

In this section we survey existing Photometric Stereo techniques. We review algorithms that improve Woodham's estimation for Lambertian surfaces as well as others that extend the estimation to non-Lambertian surfaces. We also consider colour Photometric Stereo techniques and active differential approaches. Finally we review the conditions for existence and uniqueness of the solutions.

#### Photometric Stereo for Lambertian surfaces

Significant improvements have been made on Woodham's technique by other researchers. Kulick [85] recovered a Lambertian surface using a three light Photometric Stereo technique. He proposed a scheme that minimises the error between the surface normal vectors and the recovered ones. He tested his scheme on synthetic data concluding that the reconstructed surface converged to the test surface as the quantisation noise was reduced.

Kim and Park [86] proposed a multi-image Photometric Stereo method that approximates the surface using Legendre polynomials. Their algorithm assumed

a linear reflectance and relied on brightness constraints. The performance of this Photometric Stereo approach was evaluated in terms of brightness, orientation and height.

Zhang and Shah [87] presented a scheme to estimate shape and surface albedo of a Lambertian surface using a sequence of  $n$  images. Their algorithm was formulated in the framework of a linear Kalman filter to iteratively refine the surface estimation. They concluded that the proposed algorithm showed a great improvement over classic Photometric Stereo on noisy images for both surface and albedo estimation.

### **Photometric Stereo for non-Lambertian surfaces**

Although traditional Photometric Stereo is restricted to Lambertian reflectance, subsequent approaches have extended its use to non-Lambertian surfaces. Coleman and Jain [8] proposed the use of four light sources to separate the specular and the Lambertian components in Photometric Stereo. This redundancy in the number of lights allowed them to detect and remove specularities. If one image pixel was specular, a valid surface orientation was determined with the remaining three lights. Solomon and Okeuchi [48] also proposed a four-light technique to extract the shape and surface roughness of an object that exhibits a specular lobe. They estimated shape using an algorithm similar to Coleman and Jain's. The specular pixel and the shape estimate were fitted into a simplified Torrance and Sparrow's model to determine the surface roughness.

McGunnigle and Chantler [88] proposed an algorithm that suppressed the specular reflection, so that classic Photometric Stereo could be applied. They modelled the facet reflectance as a function of the illuminant azimuth. To estimate the reflectance function for each facet they illuminated the surface from several azimuth angles. They fitted the measured reflectance into a function which had two terms: a Lambertian component and a specular residue. By suppressing the specular residue, they could apply traditional Photometric Stereo to recover the facet orientation.

Ikeuchi [89] used distributed light sources for Photometric Stereo of specular

surfaces. By using several extended sources he ensured the detection of the specular reflection. Nayar *et al.* [90] extended this approach to surfaces whose reflection is a sum of specular and diffuse components. Since the surface orientation and reflectance are estimated locally, this method can be applied to non-homogeneous surfaces. The relative strength of the Lambertian and specular components does not need to be known *a priori*, it is estimated.

Tagare and deFigueiredo [91] introduced a Photometric Stereo scheme for a large class of non-Lambertian surfaces. They proposed a reflectance model for diffuse homogeneous surfaces, called *m-lobed reflectance map*. They compared the performance of their approach to that of the standard Photometric Stereo for a non-Lambertian sphere. They found that the reconstruction from using his *m-lobed reflectance map* was superior to that from using the Lambertian reflectance map.

Kay and Caelli's method [92] simultaneously measured the surface normals and reflectance parameters. They assumed fairly general reflectance properties but employ a large number of images for the recovery. Although using a simplified Torrance-Sparrow model was sufficient for simulation, it did not give accurate reconstruction of real objects.

Among the more advanced methods are those that use local confidence measures to account for surface inter-reflections and shadowing [93] [94]. These methods use three or more images, and exploited this redundancy to detect shadows and inter-reflections.

## Colour Photometric Stereo

Colour information can improve the photometric recovery. With Lambertian reflection the colour components are dependent. Therefore, the information in a colour image (blue, green and red components) is equivalent to three grey-scale images of the same resolution. In contrast for non-Lambertian reflection, the information in the colour components supplement each other. The specular and diffuse components usually have different colour, and peaks in different directions. This helps separat-

ing specular and diffuse reflection as well as determining the surface orientation. Photometric Stereo approaches that take advantage of colour images are [95] [96] [97] [98] [99] [100].

### **Active differential approaches**

Classical Photometric Stereo techniques assume that the light source is at infinity. Differential Photometric Stereo assumes a point light source which is relatively close to the surface and the camera [101] [102] [103] [104]. The algorithm relies on the controlled motion of the light to solve for height. It only requires the solution of a linear equation and holds for a very general class of reflectance models.

### **Existence and uniqueness**

Recently, Okatani and Deguchi [105] investigated the issue of uniqueness for the three-light Photometric Stereo. They stated the conditions on the lighting configuration according to the surface reflectance for a unique solution. Previously, Tagare and deFigueiro [91] obtained constraints on the position of three light sources to ensure uniqueness of the solution. They restricted their investigation to a class of diffuse non-Lambertian surfaces.

Shape from Shading techniques recover the surface in terms of gradient. Our purpose is to recover the surface height function, therefore an integration step is required.

## **4.3 Integration techniques**

The integration of the surface normals to provide depth may be performed using either local or global approaches [106]. *Local* path integration techniques are easy to implement and computationally efficient. However, the local nature of the calculations and the propagation of errors means that accurate data is required. Furthermore, they do not exploit any assumptions about the integrability condition.

Surface integration is treated as an optimisation problem in *global* techniques. They are more robust to noise than local approaches since the surface gradient data is considered globally during the solution process.

### Local integration

Coleman and Jain’s local approach [8] was a scan algorithm that passed through all points of the image grid. Starting with an initial depth this algorithm can be used to propagate depth values according to a local approximation rule. They estimated the depth of a point  $P$  by considering the normal vectors of two points adjacent to  $P$ , computing the average tangent lines through  $P$  and interpolating.

Instead, Healey and Jain [107] considered the eight points surrounding  $P$  for more precise depth information. This method is computationally expensive and relies on the two-point method to find relative depth for irregular borders.

Wu and Li [108] exploited Green’s theorem and used multiple path-independent line integrals to compute relative depth. This method is powerful, simple to implement and valid for discontinuous surfaces. However its precision is lower than that of the two-point and eight-point methods.

### Global integration

Ikeuchi [109] presented a least squares technique for estimating surface shape from a needle map. His algorithm minimises the mean squared error between the surface derivatives and the surface orientation. Boundary conditions are necessary to generate an absolute depth map.

Horn [67] solved for both gradient and height simultaneously. He formulated the estimation problem as one of finding a surface that minimises a function which accounts for the brightness and the integrability constraints. He implemented this algorithm assuming a linear surface reflectance. He initially incorporated a departure from smoothness term in the function to minimise. This term is only necessary to prevent instability when far from the solution. It can be dropped as we get close



to the solution, so that we assure convergence to the exact reconstruction. The algorithm was applied to synthetic and real images proving to successfully recover complex wrinkled surfaces, even surfaces with discontinuities in the gradient.

Frankot and Chellappa [66] assumed that the unknown surface function satisfied the integrability condition. Furthermore, they assumed a Fourier coefficient representation for the surface and proved a theorem allowing the reconstruction in the Fourier domain. Thus an inverse Fourier transform leads to the desired data. This algorithm is explained in detail in Section 4.5.2.

Simchony *et al.* [110] presented a direct method for solving Shape from Shading whilst enforcing integrability. This method reduces to an algorithm similar to Frankot and Chellappa's when using Fourier transform and periodic boundary conditions. Since this method uses a central difference derivative operator, the estimated surface does not suffer from high frequency oscillation as in Frankot and Chellappa's case. They assessed the new algorithm by recovering a synthetically generated Lambertian sphere. Although this assessment is limited and initial conditions were necessary, the algorithm was shown to successfully recover the surface.

Petrovic *et al.* [111] enforced integrability via belief propagation across graphical models. The graphical model represented the joint probability density function of all unknown surface gradients given their initial noisy estimates and indicators whether the integrability constraint was satisfied for the whole surface. The algorithm exploited the manner in which the given function factored as a product of "local" functions, each of which depended on a subset of the random variables. Such a factorisation was visualised as a bipartite graph and was called a factor graph. This scheme utilised the surface integrability constraint to produce the maximum *a-posteriori* estimate of a valid surface. The algorithm was implemented as a Shape from Shading scheme and a Photometric Stereo technique. Good results were shown for real and synthetic images.

*Direct approaches* are global algorithms that solve directly for height [56] [60] [71]. This means that integrability problems are avoided. Also the problem of

finding a consistent surface height for the estimated gradient is avoided.

Harris [112] proposed a coupled depth/slope model for surface reconstruction. He also avoided consistency problems by simultaneously estimating height and gradient. His algorithm can handle constraints and discontinuities of any order of derivative. This approach is not strictly a global method but a surface interpolation method.

## 4.4 Discussion

In Section 4.2 and Section 4.3, we reviewed Photometric Stereo techniques and Integration techniques. We argued that Photometric Stereo is the algorithm most suited to this thesis since it can be applied to rough surfaces. We also highlighted that global integration methods are more robust to noise than local path algorithms.

In Section 4.5, we describe a benchmark shape recovery technique which combines a Photometric Stereo algorithm and a global integration method.

## 4.5 Benchmark photometric recovery technique

We present a surface recovery technique which comprises two standard components: a 3-light Photometric Stereo algorithm and a global integration method. The 3-light Photometric Stereo algorithm has been successfully used for rough surfaces estimation [59] [113] [114]. Similarly, the global integration method is well known in the literature [106] [115] [116]. This makes the proposed surface recovery technique a good choice as a benchmark technique which is used for comparison purposes in the rest of the thesis. The so-called Benchmark technique is an efficient shape estimator of rough surfaces [117] [118] [119]. Next we describe the two components of the shape estimator.

### 4.5.1 Simple Photometric Stereo scheme

This classic Photometric Stereo algorithm avoids inverse matrix calculations by placing the lights with the same zenith angle and azimuth angles in  $90^\circ$  increments [59]. It assumes Lambertian reflectance of the surface. It is expected to be inferior to the schemes discussed in the previous survey which assume more realistic reflectance models. However, in the context of this thesis the surfaces are nearly Lambertian with constant albedo. Furthermore, their main characteristic is that they are rough textures. We argued that Photometric Stereo is suitable for rough surface estimation.

For a Lambertian surface illuminated with azimuth  $\tau$  and zenith  $\sigma$ , the intensity follows the expression of equation 3.4. Taking three images of the surface under the described lighting conditions:

$$i_1(x, y) = i_0 \rho(x, y) \frac{-p(x, y) \cos \tau_1 \sin \sigma - q(x, y) \sin \tau_1 \sin \sigma + \cos \sigma}{\sqrt{p^2(x, y) + q^2(x, y) + 1}} \quad (4.2)$$

$$i_2(x, y) = i_0 \rho(x, y) \frac{-p(x, y) \cos (\tau_1 + \pi/2) \sin \sigma - q(x, y) \sin (\tau_1 + \pi/2) \sin \sigma + \cos \sigma}{\sqrt{p^2(x, y) + q^2(x, y) + 1}} \quad (4.3)$$

$$i_3(x, y) = i_0 \rho \frac{-p(x, y) \cos (\tau_1 + \pi) \sin \sigma - q(x, y) \sin (\tau_1 + \pi) \sin \sigma + \cos \sigma}{\sqrt{p^2(x, y) + q^2(x, y) + 1}} \quad (4.4)$$

Adding the equations 4.2 and 4.4 will produce a non-linear function of the surface derivatives:

$$i_{NL}(x, y) = i_1(x, y) + i_3(x, y) = \frac{2i_0 \rho(x, y) \cos \sigma}{\sqrt{p^2(x, y) + q^2(x, y) + 1}} \quad (4.5)$$

now dividing equations 4.2 and 4.3 by 4.5 we have two linear equations which are independent of albedo,  $\rho$ , and incident intensity,  $i_0$ :

$$i_p(x, y) = \frac{i_1(x, y)}{i_1(x, y) + i_3(x, y)} = \frac{-p(x, y) \cos \tau_1 \sin \sigma - q(x, y) \sin \tau_1 \sin \sigma + \cos \sigma}{2 \cos \sigma} \quad (4.6)$$

$$i_q(x, y) = \frac{i_2(x, y)}{i_1(x, y) + i_3(x, y)} = \frac{-p(x, y) \sin \tau_1 \sin \sigma - q(x, y) \cos \tau_1 \sin \sigma + \cos \sigma}{2 \cos \sigma} \quad (4.7)$$

The equation system linearly maps surface slope to image intensity, therefore the gradient can be computed as:

$$p(x, y) = \frac{1 - 2i_p(x, y)}{\tan \sigma} \sin \tau_1 + \frac{1 - 2i_q(x, y)}{\tan \sigma} \cos \tau_1 \quad (4.8)$$

$$q(x, y) = \frac{1 - 2i_p(x, y)}{\tan \sigma} \cos \tau_1 + \frac{1 - 2i_q(x, y)}{\tan \sigma} \sin \tau_1 \quad (4.9)$$

The system only needs the capture of three images at azimuth angles of  $90^\circ$  increments and the application of equations 4.8 and 4.9 to provide the estimates of the gradient field. We should bear in mind that this scheme ignores self- and cast shadows as well as inter-reflections. Nevertheless, it is fast and simple to implement. It is expected to have good behaviour for almost Lambertian surfaces of low slope angles where the effect of shadows and inter-reflections is small.

### 4.5.2 Frequency integration

We enforce integrability using Frankot and Chellappa's approach [66]. They represented the (possibly non-integrable) estimate of surface slopes with a finite set of basis functions. Then they enforced integrability by calculating the orthogonal projection of the estimated slopes onto a set of integrable slopes. Their algorithm provides a least squares fit of integrable slopes to non-integrable slopes where the integrability constraint is minimised.

We implement their algorithm using the Fourier basis functions to represent the surface. This frequency domain interpretation of the shape from gradient problem

provides a computationally efficient implementation. Consider a surface  $s$  which is continuous and has a Fourier coefficient representation

$$s(x, y) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} S(u, v) \cdot e^{-j \cdot (u \cdot x + v \cdot y)} du dv \quad (4.10)$$

where

$$S(u, v) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} s(x, y) \cdot e^{j \cdot (u \cdot x + v \cdot y)} dx dy \quad (4.11)$$

denote the Fourier coefficients of  $s$ .

The spectra of the surface derivative fields are related to the surface height spectrum by

$$P(u, v) = juS(u, v) \quad (4.12)$$

$$Q(u, v) = jvS(u, v) \quad (4.13)$$

It follows that the surface height spectrum can be calculated from the derivatives spectra

$$S(u, v) = \frac{uP(u, v) + vQ(u, v)}{j(u^2 + v^2)} \quad (4.14)$$

Equation 4.14 is not defined at the point  $(u, v) = (0, 0)$ : we cannot recover the average value of  $s$ , i.e. the absolute distance from the viewer to the surface.

Equations 4.12 and 4.13 show that the low-frequency components of the surface are attenuated in the image formation process. Thus the reconstructed surface inevitably suffers from low-frequency distortion, the severity depending on observation noise characteristics. Expression 4.14 shows that the Fourier coefficients of the estimated derivatives are combined in proportion to their frequencies. This reduces the amplification of low-frequency noise during integration.

Frankot and Chellappa implemented this algorithm as part of an iterative scheme. They enforced integrability at the cost of over-smoothing the surface estimate. We

implement Frankot and Chellappa's technique in a non-iterative manner, so that the surface is not over-smoothed. This implementation provides a least-squares fit of non-integrable slopes to integrable slopes in one pass of the algorithm.

## 4.6 Summary and discussion

In this chapter we reviewed single-image Shape from Shading techniques as well as Photometric Stereo algorithms. We identified Photometric Stereo as a possible candidate for the recovery of rough surfaces since it does not rely on smoothness constraints. We then reviewed several integration methods and concluded that global integration was the most appropriate approach.

In the last part of this chapter we described a 3-light shape recovery technique which combined Photometric Stereo estimation with global integration. Both the components of the proposed shape estimator are well-known in the literature and have been successfully used. Very little work has been carried out on the shape recovery of rough surfaces. In Chapter 5 we present two novel recovery techniques that are not restricted to smooth surfaces.

# Chapter 5

## Shape Recovery: Two Novel Methods

---

Although Shape from Shading is been extensively used to estimate the surface structure, it is been mostly applied to smooth surfaces. Recovery techniques that use less than three images rely on assumption about the shape or smoothness of the surface. Examples are the single-image Shape from Shading algorithms surveyed in Chapter 3 and most of the two-image methods reviewed in this chapter. The height recovery of rough surfaces, such as textures, has not been properly addressed before.

In Chapter 4 we presented a benchmark shape recovery method that combines a 3-light Photometric Stereo technique with frequency integration. The two components of the Benchmark technique have been recovered from the literature.

In this chapter we review 2-light Photometric Stereo methods. We discuss the conditions for a unique solution when two surface images are available. We find that by assuming a linear surface reflectance function, no smoothness constraints are needed. Rough surfaces can be recovered using only two images, thus more efficient algorithms can be proposed. Subsequently we present two novel techniques that can be applied to rough surfaces. We finally summarise the assumptions about the surface properties made for the derivation of our algorithms.

Technique	Assumptions		
	<i>shape</i>	<i>reflectance</i>	<i>other constraints</i>
<i>Onn</i>	Smoothness		Integrability
<i>Kozera</i>	Smoothness		Integrability
<i>Lee &amp; Kuo</i>	Small slopes	Linear	Triangular facet
<i>Hansson &amp; Johansson</i>	Small slopes	Linear	$\sum_i y_i = 0$
<i>Torreão &amp; Fernandes</i>	Smoothness	Linear	Close lights
<i>Yand et al.</i>	Smoothness		Convexity
<i>Caver &amp; Schalkoff</i>	Prototypical surface		
<i>Kim &amp; Park</i>		Linear	Brightness

Table 5.1: Two-light Photometric Stereo techniques

## 5.1 Two-light Photometric Stereo techniques

### 5.1.1 Review of the algorithms

In Chapter 4, we reviewed several techniques for shape estimation from shading information. We surveyed single and multiple image (Photometric) Shape from Shading methods [18] [84]. Conventional Photometric Stereo requires at least 3 images but several researchers have developed schemes that need only two images. The reported 2-light schemes are reviewed in this section.

Assuming a Lambertian surface we have to consider a system formed by two equations like expression 3.4. This system has to be solved for three unknowns: the two surface derivative maps and the albedo map. By assuming that the surface has constant albedo, the local surface normals can be determined up to a constant.

Work has been done to try to investigate the existence and uniqueness of the solutions. Table 5.1 summarises the approaches surveyed in this section and the assumptions that they make to find a unique solution. All the reviewed approaches assume a Lambertian surface with constant albedo. Other assumptions about the surface shape, its integrability or its reflectance are included in the table.



Onn [120] stated that for smooth surfaces, the local integrability constraints usually resolved the problem of deciding between the two possibilities. Kozerá [121] argued that generally this system had, up to a constant, a unique solution for a smooth Lambertian surface. He also examined exceptional cases where there was no such unique solution. He used smoothness constraints and integrability conditions to resolve the ambiguity.

Lee and Kuo [122] generalised an iterative Shape from Shading algorithm to a two-light Photometric Stereo algorithm. They assumed linear reflectance of the surface. They approximated the reflectance map  $R(p, q)$  by the first term of its Taylor series expansion about a certain reference point  $(p_0, q_0)$ . This point could be either fixed or varying for different values of  $(p, q)$ . By assuming constant albedo of the surface and a linear reflectance, they established a framework where one single solution exists for the estimation problem. Their method implied the use of a triangular-element surface model. The accuracy of the surface estimation was affected by the type of triangulation scheme and by the level of resolution chosen.

A two-light Photometric Stereo technique developed by Hansson and Johansson [116] was presented as an application for estimating paper reflectance and topography. The two images of the surface were lit from the left and right respectively. Only one of the derivative maps was estimated,  $p(x, y)$ , i.e. only the inclination of the surface elements was determined. They assumed that the mean height in the  $y$ -direction was zero. They also approximated the surface reflectance to a linear function around  $(p, q) = (0, 0)$ . This approximation was valid for small slopes. The inclinations were integrated to give a surface height function with the help of a Wiener filter, which suppressed frequencies that have a poor signal-to-noise ratio. The main characteristic of this method is that two images with opposite azimuth are enough to solve for one of the derivative fields when the reflectance is a linear function.

Torreão and Fernandes [123] [124] presented a 2-light Photometric Stereo algorithm that also assumed a linear reflectance map. When the illumination directions

of both images were not far apart, and if the imaged surface was smooth, the linear approximation to the reflectance function was applicable. The disparities produced by the matching process can be related to the depth of the imaged surface. To determine the disparity map between images, they used a constraint which is equivalent to the optical-flow-constraint equation. They also assumed a constant albedo of the surface.

Yang *et al.* [125] proposed an approach that assumed a strictly convex surface to find the unique solution. Also Caver and Schalkoff [126] solved the ambiguity by assuming a prototypical surface. The novelty of their work was that the extended light sources and the imager were close to the surface. They modelled the surface reflectance using a linear combination of three components: Lambertian, back-scatter, and extended-source Phong’s specular reflection.

Kim and Park [127] approximated the surface using Legendre polynomials and then estimated it with a two-image Photometric Stereo method. Their algorithm assumed linear reflectance and relied on brightness constraints. They compared the performance of their algorithm to that of conventional Photometric Stereo for noiseless/noisy images. The assessment was in terms of brightness error, orientation error and height error. They concluded that their algorithm gave good performance comparable to the conventional methods and that it was robust to noise.

The reviewed approaches rely on different assumptions and constraints of the surface to reach a unique solution. We next discuss the conditions for the uniqueness and existence of the solution.

### 5.1.2 Uniqueness and existence

We summarise the conditions that we use to solve for a unique surface. Most of the discussed two-light methods rely on assumptions about the smoothness or the shape of the surface. In our investigation, the restrictions of surface properties should not refer to its topography or roughness. Table 5.1 shows that by assuming a linear reflectance function, smoothness constraints can be avoided. This is going to be our

approach. A unique solution is found for a surface with:

1. Lambertian reflectance.
2. Constant albedo map.
3. A surface-to-image transfer function which is accurately modelled by a linear function.
4. A potential height function to a conservative gradient field.

Assuming a Lambertian surface we simplify the estimation problem to one of solving the a system of two non-linear equations. For a constant albedo map, only two unknowns have to be computed, i.e. the derivative fields  $p(x, y)$  and  $q(x, y)$ . Fitting the surface-to-image model to a linear function makes the equations linear, and the unique solution can be easily found up to a constant. If the surface gradient is conservative, the integrability condition is met and the gradient can be integrated into the surface.

Although this approach restricts the surface characteristics, we adopt it because the assumptions about the surface are reasonable for our textures. The Lambertian assumption is common in computer vision [128]. The reflectance function is nearly linear for a Lambertian surface with low slopes [56]. The surface has to be continuous for its gradient field to be integrable. A continuous surface consists of a succession of low slope facets and all can be seen by the viewer (from overhead). It could still be rough since two consequent facets can differ greatly in their slopes. Therefore our test surfaces are continuous Lambertian surfaces of low slope angles. The most severe restriction is assuming a constant albedo.

Traditionally, Photometric Stereo requires three or more images to accurately estimate rough surfaces. We here argued that two images are sufficient to recover rough surfaces, thus more efficient algorithms can be proposed.

## 5.2 Linear Photometric Stereo

### 5.2.1 Introduction

The Linear Photometric Stereo technique is a novel two-light method. It is based on an assumption about the surface reflectance function. This method exploits the surface-to-image function that was presented by Kube and Pentland in [55]. They presented a linear transform that relates the spectrum of the surface to the spectrum of its image (see Section 3.7. for derivation):

$$I(\omega, \theta) = -j\omega \sin \sigma \cos(\theta - \tau)S(\omega, \theta) \quad (5.1)$$

Pentland [56] developed a single image Shape from Shading technique that applies the inverse of the transform to recover the surface from the image. He considered neither smoothness constraints nor boundary conditions. The illumination conditions, azimuth ( $\tau$ ) and zenith ( $\sigma$ ), had to be known to calculate the inverse filter.

Pentland pointed out that, with a linear reflectance, the surface frequency components that are exactly perpendicular to the illuminant azimuth could not be seen in the image. He proposed that these surface frequency components must be either estimated from other information sources or set to some default value. Boundary conditions, such as surface shape from contours or singular points, can be used to compensate for the lack of information and an accurate recovery of the surface can be achieved. When boundary conditions are not available, he invokes the general view position assumption and argues that the unseen Fourier components should be set to zero. He concludes that the use of default values produces good estimates of shape for complex and irregular surfaces, but not for geometric forms. Zhang compared Pentland's algorithm to other well-known single-image Shape from Shading techniques [115]. He concluded that these methods gave poor results and suggested the use of more images to improve the estimation.

In this section we propose a technique that is designed to overcome the prob-

lem of estimating the surface polar components at right angles to the illuminant azimuth. Our approach uses two intensity images that differ in their illuminant azimuth angles. The light position is known, so the surface can be estimated from each image. The surface polar components at right angles to the light direction for the first image can be estimated from the second, and vice versa.

We expect the recovery performance of our technique to be a function of the applicability of the linear mapping of equation 5.1. This mapping is optimal for a Gaussian surface with linear reflectance. Our test surfaces are continuous with Gaussian height distributions. Furthermore, they are Lambertian surfaces with low slope angles, thus the reflectance function is nearly linear (see Section 3.8.1). However, the recovery performance may deteriorate for rougher surfaces ( $p_{rms}, q_{rms} > 0.3$ ).

### 5.2.2 Description of the algorithm

The surface can be estimated by applying the inverse of the linear filter of equation 5.1, i.e.:

$$S(\omega, \theta) = \frac{-I(\omega, \theta)}{j\omega \sin \sigma \cos(\theta - \tau)} \quad (5.2)$$

This estimation process assumes a linear reflectance. We should therefore study the circumstances under which non-linear order effects are significant.

#### Non-linear order effects

Chantler [57] investigated the validity of the linear expression (equation 5.1). He concluded that for a given surface roughness and zenith, the most interesting non-linear effects occur at  $\theta = \tau + (2k + 1)\pi/2$ . Here the linear model predicts that all components will be filtered out. Chantler showed with simulation that for the Lambertian model, quadratic and higher order terms are significant for those frequency components.

We want to maximise the ratio between the linear term and the non-linear terms of the Lambertian equation (equation 3.4). Lambert's intensity expression after

taking MacLaurin's expansion is

$$i = [-p \cos \tau \sin \sigma - q \sin \tau \sin \sigma + \cos \sigma] \left[ 1 - \frac{1}{2!}(p^2 + q^2) + \frac{9}{4!}(p^2 + q^2)^2 \dots \right] \quad (5.3)$$

Chantler predicted that the strength of non-linear effects is related to the angle between the surface frequency component ( $\theta$ ) and the illuminant azimuth ( $\tau$ ). Thus we simplify the investigation by considering a corrugated surface rotated by  $\phi$ . For a corrugated surface with  $q(x, y) = 0$  — the grain is aligned with  $y$  axis — equation 5.3 simplifies to

$$i = [-p \cos \tau \sin \sigma + \cos \sigma] \left[ 1 - \frac{1}{2}p^2 + \dots \right] \quad (5.4)$$

If we rotate the surface together with the axis system by an angle  $\phi$  without rotating the light, then equation 5.4 becomes

$$i = [-p \cos (\tau - \phi) \sin \sigma + \cos \sigma] \left[ 1 - \frac{1}{2}p^2 + \dots \right] \quad (5.5)$$

For low slope angle surfaces we can discard third and higher order terms. Rearranging equation 5.5 we get an expression with an offset term, a linear term and a second order term:

$$i = \cos \sigma - p \cos (\tau - \phi) \sin \sigma - \frac{1}{2}p^2 \cos \sigma \quad (5.6)$$

The linear term is proportional to the cosine of  $(\tau - \phi)$ , whilst the quadratic term is independent on the surface orientation and light azimuth. The linear to non-linear ratio for the corrugated surface follows:

$$\frac{S}{R} = \log \left[ \frac{\cos (\tau - \phi) \sin \sigma}{1/2p \cos \sigma} \right] = \log [\cos (\tau - \phi) \sin \sigma] - \log [1/2p \cos \sigma] \quad (5.7)$$

To improve the signal-to-residue ratio, the first logarithmic term has to be maximised and the second minimised, i.e.:




Figure 5.1: Reflectance maps of a Lambertian surface illuminated for increasing zenith and fixed azimuth.

- *The surface derivative,  $p$ , tends to zero, we are restricted to low slope angles.*
- *The illuminant zenith,  $\sigma$ , tends to  $90^\circ$ . It's worth noting that the Lambertian model neglects shadows. In reality, we have to compromise between maximising the linear term (large  $\sigma$ ) and avoiding shadows (small  $\sigma$ ). Figure 5.1 shows three reflectance maps of a Lambertian surface illuminated for increasing zenith angles and a fixed azimuth. The shallower the illumination, the more linear the reflectance map.*
- *The surface orientation,  $\phi$ , and the azimuth angle,  $\tau$ , tend to the same value. For a fixed zenith, the cosine  $\cos(\tau - \phi)$  should be maximised, i.e.  $\phi \approx \tau$ . In contrast, if  $\phi \approx \tau + \pi/2$  the cosine is close to zero and the second logarithmic term is dominant.*

Assuming that the directional components are independent, the image of the corrugated surface rotated by  $\phi$  can be used to predict the polar component of a surface image with  $\theta = \phi$ . In this sense, equation 5.6 shows that the second order term is the same for all polar components, and the linear term varies with the cosine of  $(\theta - \tau)$ .

As an example we render an isotropic texture using both Lambert's law and Kube's model. The surface is a synthetic fractal with a roll-off factor of  $\beta = 3.7$  and surface roughness  $p_{rms} = 0.25$ . The light is placed at zenith  $45^\circ$  and azimuth  $0^\circ$  for both rendering processes. Figure 5.2 shows the polar plots of the spectra of these

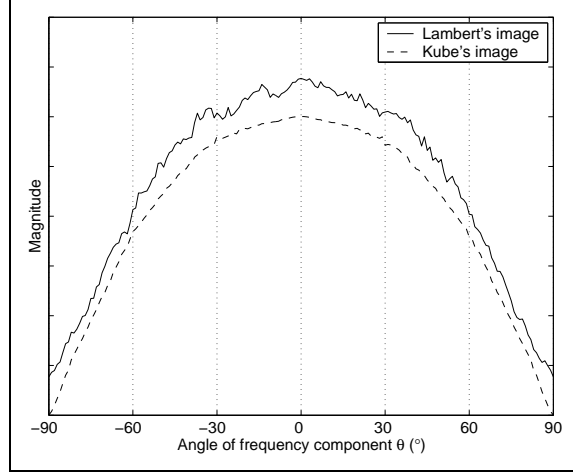


Figure 5.2: Polar plots of the spectra of two images of an isotropic fractal obtained rendering with Lambert's law and Kube's model.

two images, i.e. the sum of the power of the components along a particular radial frequency. They are different due to secondary and higher order terms. When the intensity signal is strong, the nonlinear effects can be neglected, i.e. for  $\theta \approx 0^\circ$ . The relative effect of the non-linearities is larger for  $\theta = \tau + (2k + 1)\pi/2$ , since the signal is poor i.e. for  $\theta \approx 90^\circ$  and  $\theta \approx -90^\circ$ .

The surface components with  $\theta$  in the vicinity of  $\theta = \tau + (2k + 1)\pi/2$  cannot be estimated since the image presents a poor signal-to-residue ratio.

### Considering noise

Now considering that noise is added to the intensity image, equation 5.1 takes the form:

$$I(\omega, \theta) = -j\omega \sin \sigma \cos(\theta - \tau)S(\omega, \theta) + N(\omega, \theta) \quad (5.8)$$

If we assume that the dominant noise is due to the camera, then it will be isotropic and it will be a constant term over  $\theta$  like the non-linear contribution. The Fourier components of the intensity perpendicular to the light direction are due to noise and non-linear effects. To avoid estimating the surface from the noisy data, Pentland [56] implements a Wiener filter. Noise and non-linear components of the image intensity pattern are then removed by filtering. He models the noise  $N(\omega, \theta)$  to be proportional to  $|\cos(\theta - \tau) \sin \sigma|$ , and the surface as a fractal Brownian function



whose power spectrum is proportional to  $\omega^{-4}$ .

We do not use a Wiener filter to implement our Photometric Stereo algorithm. We avoid estimating the surface from intensity components with poor signal-to-noise ratio by using two images with different illuminations.

### The algorithm

A given surface frequency component is estimated from the intensity image that shows the largest signal-to-residue ratio for that component. This is, the image whose corresponding component is most linear — closest to the azimuth direction:

$$\hat{S}(\omega, \theta) = \begin{cases} \hat{S}_1(\omega, \theta) & \text{for } [(\tau_1 + \tau_2)/2 + \pi/2] \leq \theta < [(\tau_1 + \tau_2)/2 + \pi] \\ \hat{S}_1(\omega, \theta) & \text{for } [(\tau_1 + \tau_2)/2 + 3\pi/2] \leq \theta < [(\tau_1 + \tau_2)/2] \\ \hat{S}_2(\omega, \theta) & \text{for } [(\tau_1 + \tau_2)/2] \leq \theta < [(\tau_1 + \tau_2)/2 + \pi/2] \\ \hat{S}_2(\omega, \theta) & \text{for } [(\tau_1 + \tau_2)/2 + \pi] \leq \theta < [(\tau_1 + \tau_2)/2 + 3\pi/2] \end{cases} \quad (5.9)$$

where  $\hat{S}_1$  is the surface function estimated from the image  $I_1$  using equation 5.2. We assumed that the two images differ in their zenith angle ( $\sigma_1 \neq \sigma_2$ ) and in their azimuth angle ( $\tau_1 \neq \tau_2$ ). Considering the same illuminant zenith angle for all the images simplifies the derivation of our estimator, without loss of generality. The real restriction about the light source position is in the azimuth angle which has to be different for each image. The optimal lighting set-up is investigated in Chapter 6.

### Extending the algorithm to $n$ images

Since the intensity  $I_h$  shows the best signal-to-residue ratio for values of  $\theta$  in the vicinity of  $(\tau_h + k\pi)$ , the way to combine  $n$  intensity images for the surface estimation is:




Figure 5.3: Surface recovery from 5 intensity images

$$\hat{S} = \begin{cases} \hat{S}_1 & \text{for } [(\tau_n + \tau_1)/2 + \pi/2] \leq \theta < [(\tau_1 + \tau_2)/2 + \pi] \\ \hat{S}_1 & \text{for } [(\tau_n + \tau_1)/2 + 3\pi/2] \leq \theta < [(\tau_1 + \tau_2)/2] \\ \hat{S}_2 & \text{for } [(\tau_1 + \tau_2)/2] \leq \theta < [(\tau_2 + \tau_3)/2] \\ \hat{S}_2 & \text{for } [(\tau_1 + \tau_2)/2 + \pi] \leq \theta < [(\tau_2 + \tau_3)/2 + \pi] \\ \dots & \\ \hat{S}_h & \text{for } [(\tau_{h-1} + \tau_h)/2] \leq \theta < [(\tau_h + \tau_{h+1})/2] \\ \hat{S}_h & \text{for } [(\tau_{h-1} + \tau_h)/2 + \pi] \leq \theta < [(\tau_h + \tau_{h+1})/2 + \pi] \\ \dots & \\ \hat{S}_n & \text{for } [(\tau_{n-1} + \tau_n)/2] \leq \theta < [(\tau_n + \tau_1)/2] \\ \hat{S}_n & \text{for } [(\tau_{n-1} + \tau_n)/2 + \pi] \leq \theta < [(\tau_n + \tau_1)/2 + \pi] \end{cases} \quad (5.10)$$

Again the only restriction is that the images must differ in their azimuth angles.

An example of the combination of five surface estimates obtained from five intensity images is shown in Figure 5.3.

It is an advantage that this technique can easily incorporate any number of images for the estimation. The performance is expected to improve as the number of images increases, but this will be investigated in Chapter 6.

### 5.2.3 Advantages

We review the main advantages of using the new recovery algorithm for the purpose of shape estimation. This algorithm is a simple yet effective technique, but its application is restricted to a certain class of textures (see Section 5.4).

*Advantages:*

1. The most important advantage of this method is that it can be applied to textures. No smoothness constraints are used for the estimation of the surface height.
2. Only two intensity images of the texture are necessary for the estimation. In contrast traditional Photometric Stereo requires three images.
3. The only restriction on illumination conditions is that the images must be lit from different azimuth angles. Theoretically, there is little limitation for the placement of the lights. This topic is investigated in Chapter 6.
4. To improve the estimation accuracy, any number of images can be integrated in a more efficient manner than using the Benchmark method.
5. Since we directly solve for height, we avoid the problem of finding a surface consistent with the estimated surface orientation.

Because only two images of the texture are necessary and since there are few restrictions on the lighting conditions, we argue that this is an efficient approach.

## 5.3 Optimal Linear Filter

### 5.3.1 Introduction

The Optimal Linear Filter, like the Linear Photometric Stereo, is a two-light Photometric Stereo method that assumes linear reflectance. The Linear PS algorithm is based on a theoretical prediction of the surface-to-image transfer function: the linear approximation is not optimal. In this section we describe a novel technique

that models the true reflectance function by its optimal linear approximation for a given surface. The reflectance assumptions are much weaker than for the Linear PS algorithm. An important advantage over the Benchmark 3-light recovery technique is that there is little restriction on the lighting conditions of the images. A further benefit is that we don't need to explicitly know those lighting conditions, unlike with the Linear PS method.

The Optimal Linear Filter parameters have to be trained. We estimate the linear filter that best matches the gradient-to-image function according to the least squares criterion. The training data consists of *a priori* knowledge about the surface topography as well as two surface images. An important limitation of this technique is that information about the surface is necessary for the training process.

Once the filter is calculated, it can be used to estimate the surface function. This technique is optimal for test surfaces with the same reflectance and first and second order statistics as the training surface. Furthermore it should be applied to test images with the same lighting conditions as the training image.

The filter parameters depend on the derivative operator used to compute the gradients of the training surface. Each derivative field is estimated from the potential field using a two-point estimator. The two-point estimator is the simplest form of differentiator, and it should be noted that it underestimates the magnitude of high frequencies. Nevertheless, we use it due to its simplicity.

Since we assume a linear reflectance, we are presented with a system formed by two linear equations at each pixel position. It has to be solved for two unknowns, the derivative fields at that position. Therefore, we need two optimal linear filters and two surface images for the estimation.

Some work related to the estimation of a linear surface-to-image transfer function is found in the literature. McGunnigle [59] compared the optimal linear filter with an empirically derived reflectance map. Although both the form of the filter and the training method are similar to our algorithm's, his use of the filter was different. He estimated an optimal linear filter for two synthetic surfaces to test the validity of

the linear reflectance assumption. He assessed this validity by comparing the filter’s performance in rendering terms to a reflectance map measured for a Lambertian surface. His aim was not surface recovery.

The approach described by Knill in [25] is also related to ours. In this paper an adaptive approach to solve the inverse mapping problem was proposed. Knill derived a linear estimation of surface shape from shading using a Widrow-Hoff associative learning algorithm. Our approach is more efficient since we derive an optimal closed-form solution for the surface. Knill used one single image for the recovery. This is an important difference from our work since we implement the optimal linear filter as part of a Photometric Stereo technique. Knill viewed the problem as one of statistical estimation and used Bayesian estimation techniques. He learnt the filter from a large number of synthetic fractal surfaces and their images. He then applied the filter to these and other fractals as well as to smooth surfaces. For the assessment, he compared the estimated surface normals to the synthetically generated ones. He concluded that the linear model performed well for the whole range of test surfaces. He pointed out that the accuracy of the estimation might be a function of the applicability of a linear reflectance model. This is also the case for the proposed algorithm.

### 5.3.2 Description of the algorithm

In this section we present a more detailed description of the Optimal Linear Filter implementation. We first describe the two-point derivative operator that is used to calculate the gradient field from the surface function. Then, the training process for the filter is detailed. Finally, we explain how the surface estimation process is carried out.

#### **The derivative operator**

The ideal 2D differentiator has a transfer function of the form  $H(\omega, \theta) = j\omega$ . For simplicity we use two non-ideal 1D differentiators. The two-point estimator is one of the most common derivative operators. It calculates the gradient with a

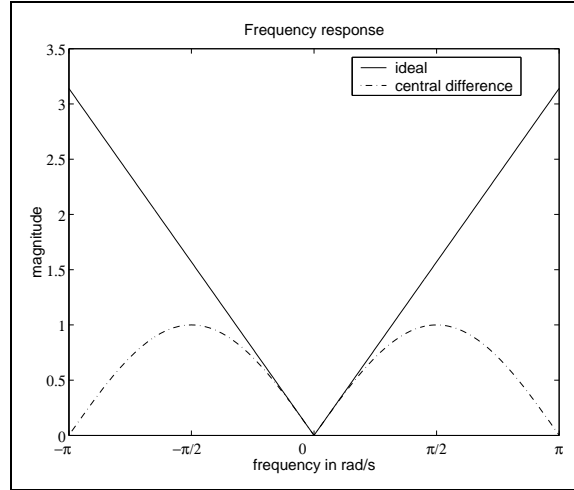


Figure 5.4: Magnitude of the frequency response for the ideal derivative operator and the central difference operator

two-neighbourhood function known as the central difference method [129]:

$$p[x_i, y_j] = \frac{1}{2}(s[x_{i+1}, y_j] - s[x_{i-1}, y_j]) \quad (5.11)$$

$$q[x_i, y_j] = \frac{1}{2}(s[x_i, y_{j+1}] - s[x_i, y_{j-1}]) \quad (5.12)$$

We compare this practical gradient estimator to the ideal case by comparing their frequency transforms. Since gradient operators are typically non-causal, we use a version of the DFT centred at zero. In general the function for the central difference method is given by:

$$h_c(n) = [\frac{1}{2}, 0, \frac{-1}{2}] \quad n = [-1, 0, 1] \quad (5.13)$$

Therefore its frequency response follows:

$$H_c(k) = \frac{1}{2} (e^{2\pi jk/N} - e^{-2\pi jk/N}) = j \sin 2\pi \frac{k}{N} \quad (5.14)$$

In Figure 5.4 we plot the magnitude of the frequency response for both the ideal differentiator and the central difference operator. The amplitude of the DFT for the central difference method follows the ideal curve quite closely for low frequencies. However the magnitude of high frequency components is underestimated.

The phase response for the ideal derivative operator is:

$$\angle H_i(\omega) = \begin{cases} -\pi/2 & \text{if } \omega < 0 \\ 0 & \text{if } \omega = 0 \\ \pi/2 & \text{if } \omega > 0 \end{cases} \quad (5.15)$$

In the case of the central difference operator, the phase response is:

$$\angle H_c(\omega) = \begin{cases} -\pi/2 & \text{if } \omega < 0 \\ 0 & \text{if } \omega = k\pi \\ \pi/2 & \text{if } \omega > 0 \end{cases} \quad (5.16)$$

The main difference between both responses is for  $\omega = k\pi$  ( $k \neq 0$ ) where the phase of the central difference operator becomes zero unlike the phase of the ideal derivative operator.

For a more detailed description of the subject see [129].

### **How to train the Optimum Linear Filter?**

Next, we tackle the problem of developing an optimum filter for a given illumination condition and also for a given surface. Applying the filter to the surface derivative fields gives an approximation to the image. The filter has the following general form:

$$\hat{i}(x, y) = \mathbf{S}^T(x, y)\mathbf{V} \quad (5.17)$$

where  $\hat{i}(x, y)$  is the image intensity at the point  $(x, y)$ ,  $\mathbf{S}(x, y)$  is the surface derivatives vector at the point  $(x, y)$  and  $\mathbf{V}$  is the filter, i.e.

$$\mathbf{S}(x, y) = \begin{bmatrix} p(x, y) \\ q(x, y) \\ 1 \end{bmatrix} \quad (5.18)$$

and

$$\mathbf{V} = \begin{bmatrix} a \\ b \\ c \end{bmatrix} \quad (5.19)$$

$p(x, y)$  and  $q(x, y)$  are the surface derivatives estimated from the surface height using two central difference operators, and  $a$ ,  $b$  and  $c$  are the filter weights.

The filter must not be linear, the surface reflectance can also be modelled with a non-linear function.

To find the optimum values for the filter parameters, the quantity  $e$  must be minimised according to a least squares criterion, where:

$$e = E[(i_d - \hat{i})^2] = E[(i_d - \mathbf{S}^T \mathbf{V})^2] \quad (5.20)$$

$e$  is the mean squared error between the desired image  $i_d$  and the approximate image  $\hat{i}$ .

The transfer function is approximated by a least squares linear filter in  $p$  and  $q$ . The filter weights only depend on the first and second order statistics of the surface, as well as on its reflectance (equations 5.21, 5.22 and 5.23). The filter is optimal for surfaces which share the reflectance characteristics and first and second order statistics with the training surface.

$$a = \frac{\overline{pq} \cdot \overline{q} \cdot \overline{i_d} - \overline{pq} \cdot \overline{i_d q} + \overline{q} \cdot \overline{p} \cdot \overline{i_d q} - \overline{p} \cdot \overline{i_d} \cdot \overline{q^2} + \overline{i_d p} \cdot \overline{q^2} - \overline{i_d p} \cdot \overline{q^2}}{2 \cdot \overline{pq} \cdot \overline{p} \cdot \overline{q} - \overline{q^2} \cdot \overline{p^2} - \overline{p^2} \cdot \overline{q^2} + \overline{p^2} \cdot \overline{q^2} - \overline{pq}} \quad (5.21)$$

$$b = \frac{\overline{pq} \cdot \overline{p} \cdot \overline{i_d} - \overline{pq} \cdot \overline{i_d p} + \overline{q} \cdot \overline{p} \cdot \overline{i_d p} - \overline{q} \cdot \overline{i_d} \cdot \overline{p^2} + \overline{i_d q} \cdot \overline{p^2} - \overline{i_d q} \cdot \overline{p^2}}{2 \cdot \overline{pq} \cdot \overline{p} \cdot \overline{q} - \overline{q^2} \cdot \overline{p^2} - \overline{p^2} \cdot \overline{q^2} + \overline{p^2} \cdot \overline{q^2} - \overline{pq}} \quad (5.22)$$

$$c = \frac{\overline{p} \cdot \overline{q^2} \cdot \overline{i_d p} + \overline{q} \cdot \overline{p^2} \cdot \overline{i_d q} - \overline{pq} \cdot \overline{p} \cdot \overline{i_d q} - \overline{pq} \cdot \overline{q} \cdot \overline{i_d p} + \overline{i_d} \cdot \overline{p^2} \cdot \overline{q^2} + \overline{i_d} \cdot \overline{pq^2}}{2 \cdot \overline{pq} \cdot \overline{p} \cdot \overline{q} - \overline{q^2} \cdot \overline{p^2} - \overline{p^2} \cdot \overline{q^2} + \overline{p^2} \cdot \overline{q^2} - \overline{pq}} \quad (5.23)$$

The surface recovery technique is a two-light Photometric Stereo technique. For each of the two images, an optimum linear filter is estimated. Therefore,



$$i_1(x, y) = \mathbf{S}^T(x, y) \mathbf{V}_1 = \begin{bmatrix} p(x, y) & q(x, y) & 1 \end{bmatrix} \begin{bmatrix} a_1 \\ b_1 \\ c_1 \end{bmatrix} \quad (5.24)$$

$$i_2(x, y) = \mathbf{S}^T(x, y) \mathbf{V}_2 = \begin{bmatrix} p(x, y) & q(x, y) & 1 \end{bmatrix} \begin{bmatrix} a_2 \\ b_2 \\ c_2 \end{bmatrix} \quad (5.25)$$

### How to estimate the surface?

Once the filter weights are determined, each filter is applied to the corresponding test image. The test image is obtained with the same lighting as the training image. The system of two linear equations in 5.26 has to be solved to estimate the gradient field.

$$\begin{cases} i_1(x, y) = a_1 p(x, y) + b_1 q(x, y) + c_1 \\ i_2(x, y) = a_2 p(x, y) + b_2 q(x, y) + c_2 \end{cases} \quad (5.26)$$

The estimated derivative fields are then integrated to recover the surface function. More than two images can be used for the estimation: an over-constrained equation system would have to be solved. The integration technique is the frequency method used for the Benchmark 3-light Photometric Stereo technique (see Section 4.5.2).

### 5.3.3 Advantages

The advantages of the Optimum Linear Filter are summarised in this section. The main drawbacks are related to assumptions about the surface properties and will be reviewed in Section 5.4.

*Advantages:*

1. The most important advantage of this method is that it can be used to recover rough surfaces.
2. Only two images are necessary to find a unique solution for the linear equation system. The solution for surface gradient exists and it is unique.

3. The illumination conditions of the two images are not restricted. Although in theory any position for the lights is possible, there are certain lighting set-ups that improve the estimation (Chapter 6).
4. The estimation performance may be a function of the applicability of a linear reflectance model. Although there is no explicit restriction on the surface reflectance and roughness, the estimator is expected to perform better for Lambertian surfaces with low slope angles.
5. Shadows are considered by the optimal linear reflectance function, i.e. this technique can predict, or at least approximate, shadows.
6. The true reflectance function is modelled by its optimal linear approximation. This is an advantage over the proposed Linear Photometric Stereo technique which used a linear approximation to a parametric model of the reflectance function.

Since information about the surface has to be known before the estimation, the usage of this technique is limited. On the other hand, it is not difficult to imagine an application where the extra information about the test data is available. For instance, this method could be implemented as part of a supervised classifier.

## 5.4 Assumptions

This section presents a summary of the assumptions that were made for the derivation of the proposed algorithms. We consider the two novel techniques as well as the Benchmark 3-light Photometric Stereo described in Chapter 4. Table 5.2 includes a short review of the surface topographic and reflectance characteristics assumed by each recovery technique.

In terms of *surface roughness*, both the Linear Photometric Stereo technique (LPS) and the Optimal Linear Filter (OLF) can only be applied to low slope surfaces.

The 3-light Photometric Stereo technique (3PS) makes no assumptions about

	3PS	LPS	OLF
<i>Surface roughness</i>	no restriction	low slope angles	low slope angles
<i>Shadows</i>	neglected	neglected	modelled
<i>Surface reflectance</i>	Lambertian	linear Lambertian	linear
<i>Albedo</i>	any	constant	constant
<i>Integrability</i>	continuous surface	continuous surface	continuous surface

Table 5.2: Summary of the assumptions made during the derivation of the three considered techniques

the roughness of the texture. On the other hand, *shadows* increase with surface roughness. They are neglected during the estimation. The roughness of the texture is therefore restricted to avoid shadowing.

Only the Optimal Linear Filter adapts shadows. If the filter is trained on shadowed data, the linear approximation to the true reflectance will account for the shadowing. Although this does not seem to be the case in practice (see Chapter 7).

In terms of *surface reflectance*, the surface is expected to be linearised Lambertian for the Linear Photometric Stereo technique. The 3-light Photometric Stereo technique assumes a perfectly Lambertian surface, and the Optimal Linear Filter assumes a surface with a linear reflectance function.

For the two-image surface recovery techniques, we assumed a constant *albedo* surface. For the 3-light Photometric Stereo, we used a third image to solve for a unique solution. No assumptions about the surface albedo were made.

To avoid *integrability* problems, we assumed that the surface gradient was integrable into the surface. The surface has to be a potential function to a conservative gradient field. We restrict the data set to continuous surface functions.

## 5.5 Summary and discussion

This chapter presented our contribution to two-image textured surface recovery. We started with a survey on two-light Photometric Stereo techniques. Some of the re-

viewed techniques rely on smoothness assumptions which limits their application to rough surfaces. We then presented two novel two-light Photometric Stereo techniques. They relied on assumptions about the surface reflectance. Their restrictions on surface topographic properties were less severe, thus they can recover textures. The assumptions made to derive our techniques represent limitations for their application, thus they can be thought of as shortcomings. An assessment to study the scope of those shortcomings is necessary. This assessment is carried out in Chapter 7, Chapter 8 and Chapter 9 on synthetic and real data.

Chapter 6 is devoted to optimise the conditions for the implementation of the new recovery methods. The optimal number of images as well as the best lighting geometry are investigated.

# Chapter 6

## Optimal Lighting Conditions

---

In Chapter 5 we presented two novel surface estimators: the Linear Photometric Stereo technique and the Optimal Linear Filter. We want to maximise the accuracy of recovery and minimise the number of images. The accuracy of the algorithms is affected by the direction from which the surface is lit. The optimal number of light sources as well as the best position for them are investigated here. The aim of this chapter is to find the best lighting geometry. This will allow us to implement the proposed recovery methods in an optimal way.

Most of the conclusions drawn here are specific to the estimation techniques. Whenever the work can be generalised to any surface recovery method it will be stated. In contrast, the investigation is not specific for a particular class of textures, it is general. We justify our assertions both by reviewing the literature and through the use of simulation. Simulation provides the necessary ground truth data and simplifies the assessment.

## 6.1 Chapter organisation

First, we investigate the optimal position for the light sources. In Chapter 5, we argued that only two images were necessary for the recovery using the proposed estimators. We first concentrate on finding the best source position for these two images. Then we extend our investigation to  $n$  images.

In the second part of our assessment we search for the optimal number of images. Adding more lights gives a more robust estimation at the cost of having a more expensive scheme. We try to find a trade-off between the accuracy and the cost of the estimation.

We summarise the optimal lighting set-up for both of the proposed techniques. The conditions for the most efficient implementation of our techniques are also stated.

## 6.2 Optimal light source position

### 6.2.1 What is the optimal position for two lights?

To answer this question we first survey the literature. We find an optimal lighting set-up for two images. Then we use simulation to assess if the proposed set-up is also optimal for our recovery methods.

#### Lee and Kuo's investigation

Lee and Kuo [122] proposed a two-light Photometric Stereo scheme and examined the optimal illumination condition that leads to the best shape reconstruction. They investigated the best lighting condition in terms of the reflectance-map function, thus their study can be generalised to any Photometric Stereo estimation.

Note that the gradient of our surface facets is expected to be concentrated in the centre region of the gradient space, say  $-0.5 < p, q < 0.5$ . Thus, this investigation is restricted to this region.

For a Lambertian surface, the image-irradiance equation is defined by:

$$E(x, y) = R(p(x, y), q(x, y)) \quad (6.1)$$

where  $E(x, y)$  is the image irradiance and  $R(p, q)$  is the reflectance map. To better understand the problem the characteristics of the reflectance map are studied. For simplicity, let's fix the value of  $q = q_o$  and view the reflectance map  $R(p, q_o)$  as a function of one variable,  $p$ . The sensitivity of  $p$  with respect to the change in  $E$  can be estimated by means of:

$$\left| \frac{\Delta p}{\Delta E} \right| = \left| \left[ \frac{\partial R(p, q_o)}{\partial p} \right]^{-1} \right| \quad (6.2)$$

This is inversely proportional to the slope of the reflectance map at  $p$ . Thus, for a fixed value of  $\Delta E$ , the value of  $\Delta p$  is smaller for the region where  $R(p, q_o)$  is steepest, i.e. where the function changes most rapidly. Therefore, the estimate  $\hat{p}$  is most accurate for the region where  $R(p, q_o)$  is steepest. Similar arguments can be made along the  $q$  direction, i.e.

$$\left| \frac{\Delta q}{\Delta E} \right| = \left| \left[ \frac{\partial R(p_o, q)}{\partial q} \right]^{-1} \right| \quad (6.3)$$

For a given point  $(p_o, q_o)$ , the sensitivity defined in equations 6.2 and 6.3 is greatest along the direction perpendicular to the contour, i.e.

$$\nabla R(p_o, q_o) = [R_p(p_o, q_o), R_q(p_o, q_o)] \quad (6.4)$$

and lowest along the tangential direction, i.e.  $[-R_q(p_o, q_o), R_p(p_o, q_o)]$ .

The estimator is provided with two images of the scene, lit from different directions and with distinct reflectance maps. It is desirable to incorporate reflectance maps that compensate for one another's weaknesses. The effect of the light position on the reflectance map has to be studied.

The unit source vector takes the form [122]:

$$L = \frac{(-p_s, -q_s, 1)^T}{\sqrt{1 + p_s^2 + q_s^2}} = (\cos \tau \sin \sigma, \sin \tau \sin \sigma, \cos \sigma)^T \quad (6.5)$$

This is,

$$p_s = -\frac{\cos \tau \sin \sigma}{\cos \sigma} \quad (6.6)$$

and

$$q_s = -\frac{\sin \tau \sin \sigma}{\cos \sigma} \quad (6.7)$$

In the gradient space, the angle  $\chi$  that the line passing through  $(0, 0)$  and  $(p_s, q_s)$  makes with the p-axis is

$$\chi = \arctan \left( \frac{q_s}{p_s} \right) = \arctan (\tan \tau) = \tau \quad (6.8)$$

and the distance  $d$  between  $(0, 0)$  and  $(p_s, q_s)$  is

$$d = \sqrt{p_s^2 + q_s^2} = \tan \sigma. \quad (6.9)$$


Note that the angle  $\chi$  is exactly the same as the azimuth angle, and the distance  $d$  depends on the zenith angle. The azimuth angle determines the orientation of the reflectance map around the origin (compare Figure 6.1(a) and (b)). The zenith angle determines the shape of the reflectance map (compare Figure 6.1(a) and (c)).

Lee and Kuo state that for a zenith angle in the range  $30^\circ$  to  $60^\circ$ , the reflectance map covers the central region of the gradient space ( $-0.5 < p, q < 0.5$ ). Thus they concentrate on finding the optimal lighting condition in terms of the azimuth angles of the two light sources. The difference of orientation between two reflectance maps with azimuth angles  $\tau_1$  and  $\tau_2$  is simply  $|\tau_1 - \tau_2|$ . It was argued that the reflectance map gives good sensitivity along the direction perpendicular to the contour but poor sensitivity along the tangential direction. We want the *gradient direction*<sup>1</sup> of one reflectance map to correspond to the tangential direction of the other reflectance


---

<sup>1</sup>Direction of maximum rate of change







(a)  $(\tau, \sigma) = (45^\circ, 45^\circ)$



(b)  $(\tau, \sigma) = (135^\circ, 45^\circ)$



(c)  $(\tau, \sigma) = (45^\circ, 60^\circ)$



(d) Combined reflectance maps (a) and (b)

Figure 6.1: Contour plots of Lambertian reflectance maps

map over the region of interest. This can be achieved by choosing the difference between the azimuth angles to be  $90^\circ$ . One such example is shown in Figure 6.1(d), where the contour plots of two perpendicular reflectance maps are shown together.

Lee and Kuo also presented some experimental work on the subject [122]. They kept one light stationary and varied the azimuth of the second. They captured images of the surface at  $15^\circ$  increments of azimuth for the second light. The rms error of the reconstructed surface heights was tabulated against the azimuth angle. These experimental results confirm that the best lighting conditions occur for a  $90^\circ$  increment between the azimuth angles of the lights. Although this investigation is for a particular surface recovery method, it still gives an insight into the best illumination conditions.

## Discussion

After reviewing Lee and Kuo's best lighting conditions, we study how this relates to our situation. Since our estimators assume a linear reflectance function, we should avoid lighting conditions for which non-linear effects are large.

To find the best zenith angle, we should consider shadowing effects as well as non-linear effects. For raking light the linear approximation to the surface reflectance is more accurate, but shadows increase (see Section 5.2.2). In contrast, for small zenith angles non-linear effects are dominant. A range of zenith angles that cover the mid region seems a good compromise, thus we restrict the lighting to  $30^\circ < \sigma < 60^\circ$ .

We recall that the accuracy of the linear approximation for unidirectional surfaces varied with azimuth (see Section 3.8.1). Non-linear terms are dominant when an anisotropic surface is lit from a direction parallel to the surface grain [59]. Although this effect should be considered, we start our investigation by adopting Lee and Kuo's conclusion for the optimal light azimuths.

To summarise, the optimal lighting conditions can be written as:

$$\Delta\tau = |\tau_1 - \tau_2| = 90^\circ \text{ and } \sigma \in [30^\circ, 60^\circ]$$

## Validation through simulation

Simulation is used to find the best lighting conditions for two lights. Two different experiments are carried out. In the first, we test the predictions of optimal azimuth positions. In the second experiment, we assess the effect of varying the azimuth angle on the recovery of very directional surfaces.

### Is $\Delta\tau = 90^\circ$ optimal?

**Method** The relative position between the azimuth angles of the two lights is varied. The zenith is fixed to  $45^\circ$  throughout the experiment. The surface is then estimated using both methods.

We do not consider the position of the lights relative to the surface directionality. Thus we only use isotropic surfaces. The test textures are the fractal and the Mulvaney surfaces. Their roughness is set to  $p_{rms} = q_{rms} = 0.3$  and their reflectance follows Lambert's law.

In this assessment, we fix one of the lights to azimuth  $\tau_1 = 0^\circ$ , whilst we vary the second light's azimuth in  $10^\circ$  increments within the range  $50^\circ \leq \tau_2 \leq 130^\circ$ . Thus the recovery performance is assessed for a range of illuminations such as  $50^\circ \leq \Delta\tau \leq 130^\circ$ . This is,  $\Delta\tau$  is varied around the value that was predicted to be the optimum.

The quantity used to measure the estimation accuracy is the signal-to-residue ratio defined by:

$$\frac{S}{R} = 10 \log \left( \frac{\text{var}[s]}{\text{var}[s - \hat{s}]} \right) \quad (6.10)$$

where  $s$  is the generated surface function,  $\hat{s}$  is the estimated surface function and  $\text{var}[x]$  is the variance of the process  $x$ . For ease of discussion we set a threshold of 10 dB for an accurate estimate.

The  $S/R$  ratio is measured and averaged across the fractal and the Mulvaney estimations. Then, the averaged  $S/R$  ratio is plotted for both techniques against  $\Delta\tau$ .




Figure 6.2: LPS estimation accuracy against the position of the lights for isotropic synthetic textures.




Figure 6.3: OLF estimation accuracy against the position of the lights for isotropic synthetic textures.

**Results and discussion** The effect of varying  $\Delta\tau$  is shown for the Linear Photometric Stereo technique in Figure 6.2 and for the Optimal Linear Filter in Figure 6.3.

As predicted the surface is estimated most accurately when the lighting directions are perpendicular to each other. The surface recovery performance is very sensitive to the relative position of the lights. The accuracy of the Optimal Linear Filter drops off rapidly compared to that of the Linear Photometric Stereo. Although it is accurate for the range  $80^\circ < \Delta\tau < 100^\circ$ . This restricts its use to lighting setups where both lights are separated nearly  $90^\circ$ . For the Linear Photometric Stereo technique, the range of possible lighting directions is wider,  $65^\circ < \Delta\tau < 120^\circ$ . It does not require a set of near perpendicular lights for an accurate estimation.

## What is the optimal lighting direction for anisotropic surfaces?

**Method** The performance of the proposed methods deteriorates when non-linear effects dominate the reflectance function. This happens when a very anisotropic surface is lit from a direction parallel to the surface grain. Next we assess the robustness of our techniques against this effect.

The test surfaces should be strongly anisotropic, so that they can be considered unidirectional. We use an Ogilvy surface with cut-off frequencies of 32 cpi and 16 cpi for the horizontal and vertical directions respectively. The directionality of the surface is set to  $d = 0.57$  to ensure unidirectionality. The surface roughness is set to  $p_{rms} = 0.30$  and  $q_{rms} = 0.17$  and the reflectance is Lambertian. The surface grain is aligned to the  $y$  axis ( $90^\circ$  direction).

In this experiment we fix the relative position between lights to the value that was predicted and measured to be optimum, i.e.  $\Delta\tau = 90^\circ$ . We simultaneously move both lights with respect to the surface in  $5^\circ$  azimuth increments. The first light's azimuth is varied from  $0^\circ$  to  $90^\circ$ ,  $0^\circ \leq \tau_1 \leq 90^\circ$ . Thus the second light is moved in the range  $90^\circ \leq \tau_2 \leq 180^\circ$ . The zenith angles of the lights are fixed and equal,  $\sigma_1 = \sigma_2 = 45^\circ$ .

The performance is measured for the two surface recovery methods in terms of signal-to-residue ratio (equation 6.10). The results are plotted against  $\tau_1$ .

**Results and discussion** Figure 6.4 shows the recovery performance of the Linear Photometric Stereo technique. We measure an increase in accuracy of estimation for  $\tau_1 = 45^\circ$  and  $\tau_2 = 135^\circ$ . This is when the distance between both the lights' azimuths and the grain direction is maximum. When  $\tau_1 = 0^\circ$  ( $\tau_2 = 90^\circ$ ) then the second light is parallel to the grain and the estimation deteriorates. Similarly, for  $\tau_1 = 90^\circ$  ( $\tau_2 = 180^\circ$ ), the first light is parallel to the surface directionality and the estimation loses accuracy. When one of the lights is aligned to the surface grain, the SRR drops by almost 1 dB. The effect is noticeable but it is not as strong as the measured effect for relative position between lights (previous section). Although the




Figure 6.4: LPS estimation accuracy against the position of the lights relative to the surface grain of an unidirectional synthetic texture




Figure 6.5: OLF estimation accuracy against the position of the lights relative to the surface grain of an unidirectional synthetic texture

estimation is not very sensitive to this effect, it is not convenient to lit an anisotropic surface from a direction parallel to its grain.

Figure 6.5 shows the sensitivity of the Optimal Linear Filter to the light position. Again when one of the two lights is parallel to the surface grain, i.e.  $\tau_1 = 0^\circ$  or  $\tau_1 = 90^\circ$ , the estimation deteriorates by almost 0.5 dB. When neither of the light azimuths is aligned with the surface directionality ( $\tau_1 = 45^\circ$  and  $\tau_2 = 135^\circ$ ), the SRR is largest. The Optimal Linear Filter performance is quite robust to this effect.

We conclude that the difference in the direction of lights has a larger effect on surface estimation than the position of lights relative to the surface. As long as the two lights are  $90^\circ$  apart, the surface estimation can be assumed to be robust to

surface directionality.

### 6.2.2 What are the best azimuths for $n$ lights?

For two lights, estimating the surface with the Optimal Linear Filter means solving a system of two linear equations for two unknowns. When using more than two images the set of equations is over-constrained. We could solve the over-constrained system using for instance a least squares technique. Instead we choose a simpler approach where we average the surface estimates. This means that the constrained system is solved for each pair of images to give  $n/2$  estimates. Then the surface estimates are averaged. The Optimal Linear Filter is always applied to two images. Since, the best azimuth positions for these two images were found to be  $\Delta\tau = 90^\circ$ , they don't need further investigation.

In contrast, the Linear Photometric Stereo technique can readily incorporate more than two images into the recovery process. In Section 5.2.2, we saw how to combine the information of several images to obtain more accurate estimates of the surface. By controlling the lighting conditions of the images, the estimation can be further improved.

Kube predicts that the intensity dependency on the zenith angle is constant and equal to  $\sin \sigma$  [55]. We choose the same zenith angle for the  $n$  images, i.e. they are scaled by the same constant. Zenith must be within the range  $30^\circ \leq \sigma \leq 60^\circ$ , here we set it to  $45^\circ$ .

The surface information contained in two images with the same zenith and opposite azimuths is approximately the same. This approximation is only valid in the absence of shadows. Since our test textures have low slope angles, shadows are small and the approximation is accurate. Furthermore, shadows are neglected by our surface estimator, so this approximation does not imply extra assumptions. We can limit the investigation into the optimal light position to azimuth angles in the range  $0^\circ \leq \tau < 180^\circ$ .

The surface reflectance is assumed to be linear, so we search for the light position

that makes the linear term dominant. The image spectrum is near linear for those polar components that are close to the light azimuth direction, i.e. when  $\theta$  is in the vicinity of  $\tau$  (see Section 5.2.2). We do not assume any particular distribution of the facet azimuths. If we place the lights evenly distributed over the considered azimuth range, the near linear intensity components are evenly distributed over the polar angle. The rule to assign the illuminant azimuth angles to  $n$  images is:

$$\left\{ \begin{array}{l} \tau_1 = \tau \\ \tau_2 = \tau + \pi/n + (2k+1)\pi \\ \tau_3 = \tau + 2\pi/n + (2k+1)\pi \\ \dots \\ \tau_i = \tau + (i-1)\pi/n + (2k+1)\pi \\ \dots \\ \tau_n = \tau + (n-1)\pi/n + (2k+1)\pi \end{array} \right. \quad (6.11)$$

Each component of the surface spectrum is estimated from the corresponding component of the image spectrum. Since we have  $n$  images, for each surface component we have  $n$  intensity components. We estimate the surface using the intensity component that is best approximated by the linear term, i.e. the one that is closest to the light azimuth direction.

When two images are used for surface estimation the even sampling condition implies that the difference between both azimuth angles is  $90^\circ$ . We drew the same conclusion when the optimal azimuth angles for two lights were investigated.

In Figure 6.6(a) we plot the approximately linear reflectance function against  $(\theta - \tau)$ . For  $(\theta - \tau) = 0, \pi$  and  $2\pi$ , the linear term is large. Noise and non-linearities are expected to be small compared to the linear term (see Section 5.2.2). Thus the linear approximation is valid. For  $(\theta - \tau) = \pi/2$  and  $3\pi/2$ , the signal vanishes. At these angles, we expect non-linearities and noise to be larger than the signal, so the linear approximation is not valid (see Figure 5.2). Figure 6.6(b) shows how two images with evenly sampled illumination ( $\tau_1 = 0^\circ$  and  $\tau_2 = 90^\circ$ )



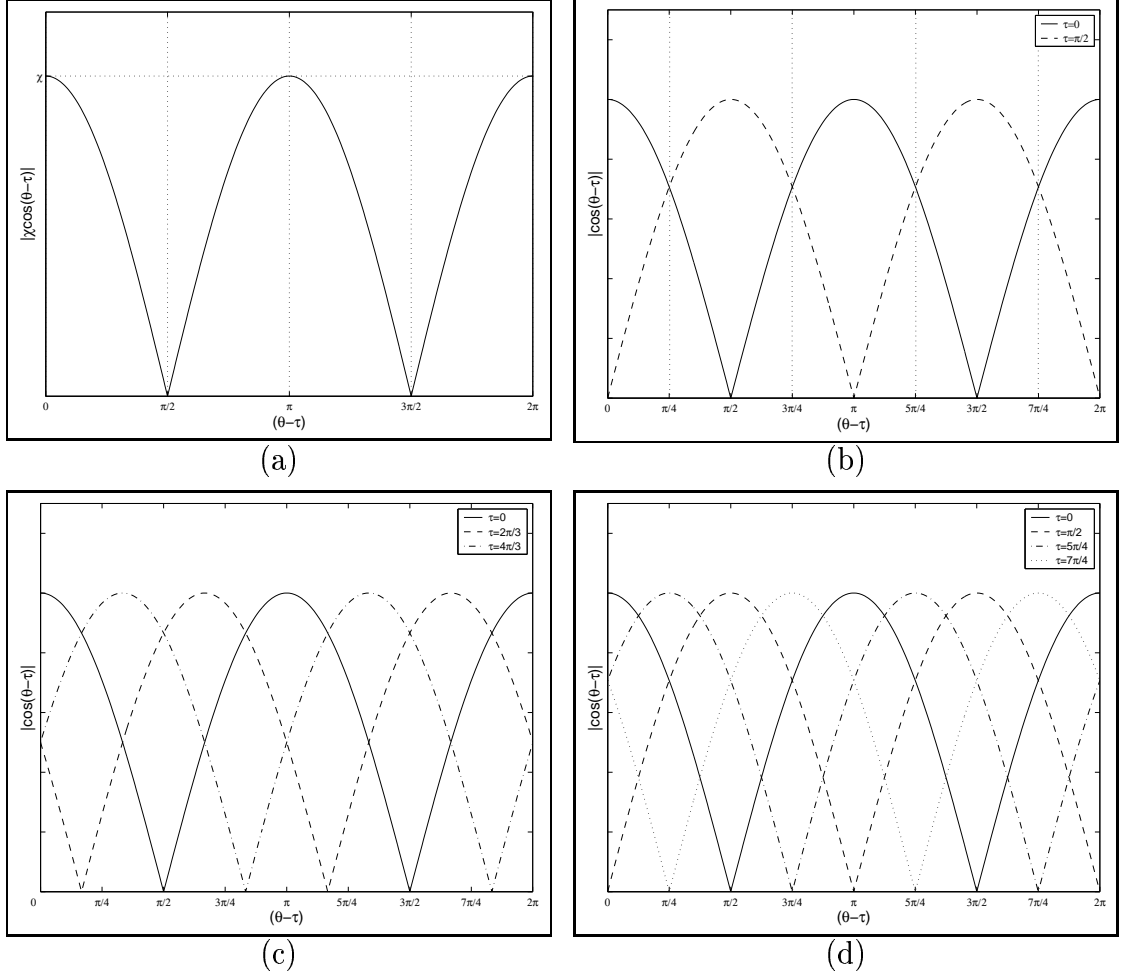


Figure 6.6: Dependency of intensity with  $(\theta - \tau)$  for (a) one image, (b) two images, (c) three images and (d) four images. The lights are evenly placed within the range  $0^\circ \leq \tau \leq 180^\circ$ .

compensate for one another's weaknesses by presenting the near linear components in different parts of the spectrum. In Figure 6.6(c) and (d) the intensity is plotted against  $(\theta - \tau)$  for three and four images whose azimuth angles are evenly placed throughout the considered range. Similarly, the near linear intensity components are evenly distributed throughout the polar spectrum, so that the surface can always be accurately estimated from one of the images.

## 6.3 Trade-off on the number of lights

### 6.3.1 Introduction

There is no restriction on the amount of images that we can use to estimate shape. The larger the number of images the more accurate the estimation but the more expensive the recovery. Here we use simulation to search for a compromise.

In theory two images are enough to find a unique solution for the surface. But, is this solution accurate? Is it necessary to improve the estimation by using more data?

Most surface recovery techniques only use three lights. The Benchmark 3-light Photometric Stereo uses three images for the estimation under no assumptions about surface albedo or the linearity of the reflectance function. It is therefore more expensive than our estimators, but it can be applied to a wider range of problems. Using four or more images implies an expensive estimation, although it may be justified depending on the application.

### 6.3.2 Simulation

#### Experimental method

We described how to integrate the information from  $n$  images when using the Linear Photometric Stereo (equation 5.10). Similarly, we stated that for the Optimal Linear Filter the  $n$  images were combined in groups of two images. The surface was estimated for each pair of images and the estimates were then averaged.

In this experiment, several synthetic surfaces are recovered with both methods using two, three and four images per texture. The generated surfaces consist of three fractal Brownian textures with surface roughness of  $p_{rms} = 0.1, 0.2$  and  $0.3$ . We assume Lambertian surfaces. The signal-to-residue ratio defined in equation 6.10 is measured for each estimate. Then the measured SRR are averaged across the three generated surface types for a more robust assessment.

<i>Technique</i>	<i>2 images</i>	<i>3 images</i>	<i>4 images</i>
<i>LPS</i>	28.96 dB	29.37 dB	29.68 dB
<i>OLF</i>	28.84 dB	30.35 dB	34.20 dB

Table 6.1: Recovery accuracy against number of images

The surfaces are illuminated under the most appropriate conditions for each recovery. We fix the illuminant zenith angle to  $\sigma = 45^\circ$  throughout the experiment. For the Linear Photometric Stereo estimator, the best illumination conditions occur for an even sampling. For the Optimal Linear Filter, the most accurate estimation occurs for light azimuths  $90^\circ$  apart.

## Results and discussion

The averaged SRR is presented in the Table 6.1 for both techniques.

Using more than two lights slightly improves the estimation for both methods. It also makes the schemes more expensive. In our case, the performance improvement does not justify the use of more than two lights for either estimation. It is for the reader to decide if the use of extra images is justified depending on the application.

## 6.4 Summary and discussion

In this chapter we considered the optimal lighting geometry. We investigated the optimal position of lights for the two novel estimators. We reviewed the literature and conducted simulation. We also searched for a compromise between number of images and accuracy of estimation. We concluded that increasing the number of lights gave a small improvement. This improvement was too small to justify the extra expense — in most foreseeable applications. The optimal lighting conditions for two lights were found to be

- Zenith angles within the range  $30^\circ \leq \sigma < 60^\circ$ .
- Azimuth angles with  $90^\circ$  increment.

These conclusions should be considered for the implementation of the novel techniques.

Once the optimal implementation for our methods is found, we are ready to assess their recovery performance. For the assessment, we should consider the limitations that were identified in Chapter 5. Therefore, the aim of subsequent chapters is to assess the algorithms on synthetic and real data.

# Chapter 7

## Assessment of Surface Estimation: Simulation

---

In Chapter 4 we described a benchmark three-light Photometric Stereo approach (3PS) from the literature [59]. In Chapter 5 we proposed two novel schemes that require fewer images, the Linear Photometric Stereo (LPS) and the Optimal Linear Filter (OLF). We also reviewed several assumptions made by our estimators and argued that they could be interpreted as shortcomings. In this chapter we assess the robustness of our algorithms to the effect of four experimental conditions related to the reviewed shortcomings:

- The type and roughness of the test surface.
- Surface reflectance.
- Shadowing.
- Noise.

The work described in this chapter allows us to state the expected accuracy and scope of the algorithms.

We use simulation to assess the performance of the algorithms. Simulation is never as convincing as experiments that use real data. However it has two critical advantages for this chapter. First, it provides a ground truthed data set. Thus the accuracy of the recovery can be measured. Secondly, it allows the experimental conditions to be varied in a controlled manner. Thus the effects of these conditions can be studied individually, and in detail, allowing a more analytical approach. By using simulation we are able to predict the accuracy of the algorithms under specific experimental conditions and to identify the reasons for that behaviour.

## 7.1 Simulation Framework

All the simulations in this chapter use the same framework. The framework has four parts which correspond to the experimental conditions:

**Surface topography model** . Three models from the rough surface literature are used: a Brownian fractal model [30] and models proposed by Mulvaney [31] and Ogilvy [23]

**Reflectance function model (local)** . Phong's model is used to predict the appearance of the surface. Ignoring the term of ambient reflection, the model has three variables,  $k_d$  controls the strength of the Lambertian component,  $k_s$  controls the strength of the specular component, and  $n$  controls the width of the specular lobe. If  $k_s$  is set to zero, the model obeys Lambert's law.

**Shadowing (global)** . A simple model for self and cast shadowing is used: a facet that is found to be shadowed is given an intensity of zero.

**Noise** . We model noise as a random process that is additive, Gaussian, white and independent of the signal. These assumptions are common in the literature [130]. Different realisations of the same random process are added to the photometric images. The amount of noise is specified with respect to the image as a signal to noise ratio (SNR).

Surface roughness, surface reflectance and shadowing effects are not independent variables. Thus they should not be considered separately. The effect of increasing the surface roughness has been considered together with shadowing effects. However, the reflectance characteristics of the surface are assumed to be independent of its roughness. This approximation is valid for Lambertian surfaces with low slope facets.

The algorithms under test can be applied to the simulated images. Throughout this chapter we will assess the accuracy of surface estimation in terms of the signal to residue ratio (SRR), equation 7.1. For ease of discussion we set a nominal lower bound of 10dB for an 'accurate' estimate of the surface.

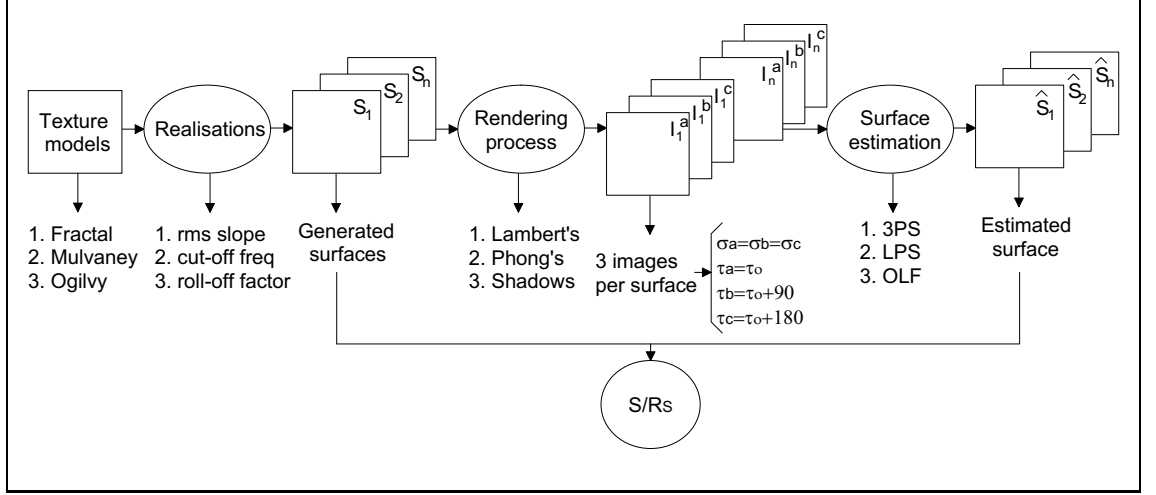


Figure 7.1: General scheme for the surface estimation assessment

$$\frac{S}{R_s} = 10 \log \left[ \frac{\text{var}[s(x, y)]}{\text{var}[s(x, y) - \hat{s}(x, y)]} \right] \quad (7.1)$$

where  $s(x, y)$  is the generated height function,  $\hat{s}(x, y)$  is the estimated height function and  $\text{var}[x]$  is the variance of  $x$ .

The assessment details are described in the scheme of Figure 7.1

## 7.2 Effect of Surface Roughness

### 7.2.1 Motivation

This section investigates the sensitivity of the proposed techniques to changes in surface roughness. Our linear recovery techniques assume an approximate reflectance which is valid for fairly smooth surfaces. Kube's approximation is reasonable where the slope angles of the surface are low  $p, q < 0.3$  (see Section 3.8.1). For a rougher surface, we expect the performance of Linear Photometric Stereo to deteriorate. Furthermore, the standard deviation of the Gaussian slope distribution is large. Thus the least squares fitting to the surface-to-image process is less accurate and the Optimal Linear Filter estimation deteriorates.

Self and cast shadows are modelled and added to the images for a more realistic rendering. Here we do not study the effects of shadowing, that is the aim of Section

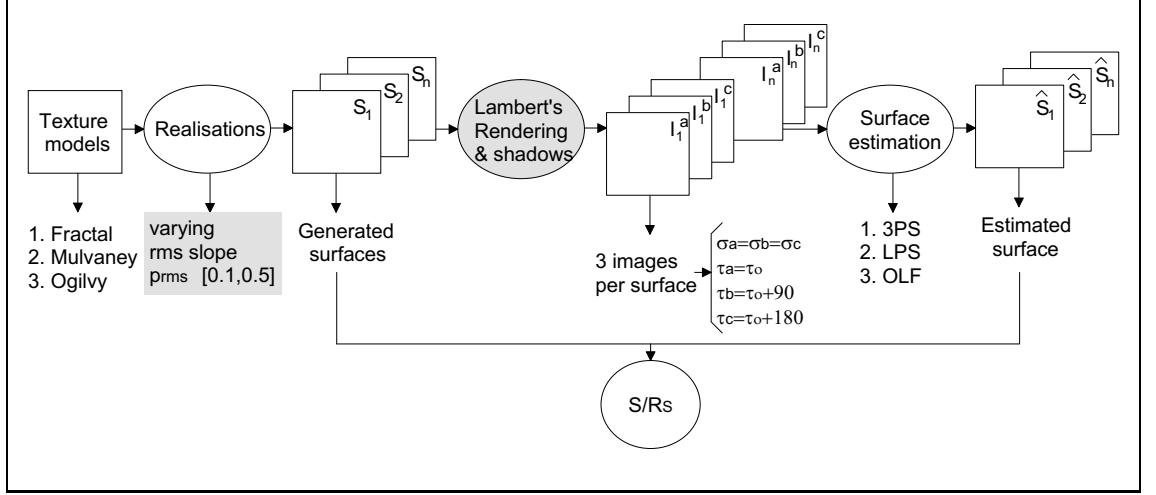


Figure 7.2: Procedure for assessing surface estimation against surface roughness

7.3. The Linear PS and the 3-light PS neglect shadowing during the estimation process. Because the amount of shadowing increase with surface roughness, we expect these estimations to deteriorate for a rougher surface.

With this simulation, we measure the maximum surface roughness for which our estimates are accurate.

## 7.2.2 Method

We carry out an experiment where the surface roughness is systematically increased. The generated textures are then rendered with Lambert's law and shadows are modelled. The three proposed algorithms are fed with these data and the estimation accuracy is measured.

Surface roughness is parameterised by the rms slope. The rms slope of the isotropic textures is varied from 0.1 to 0.5 in 0.05 increments. For the anisotropic Ogilvy,  $p_{rms} \in [0.1, 0.5]$  in 0.05 increments, whilst  $q_{rms} \in [0.076, 0.385]$  with an increment of 0.039.

The  $S/R_s$  ratio defined as in equation 7.1 is measured for each texture. Then, we average the  $S/R_s$  across the three textures for each roughness. Figure 7.2 summarises the experimental approach.






Figure 7.3: Surface recovery performance against rms slope for synthetic Lambertian textures (including shadows)

### 7.2.3 Results and discussion

The  $S/R_s$  ratio for the three surface estimates is plotted against roughness in Figure 7.3.

The Benchmark 3-light Photometric Stereo performance (3PS) is superior over the whole roughness range. The Lambertian rendering process is fairly well inverted giving an accuracy of 20 dB for low slope angles ( $p_{rms}, q_{rms} < 0.3$ ). As the slopes get steeper, shadows cause the accuracy of the 3-light PS to drop by 5 dB. However, the signal-to-residue ratio is over 10dB for the considered roughness range.

The performance of both the Linear Photometric Stereo technique (LPS) and the Optimal Linear Filter (OLF) also deteriorate for increasing roughness. The estimation accuracy is acceptable for rms slope smaller than 0.35. The causes for the deterioration with increasing roughness are: a less valid linear approximation and an increase in shadowing. We should note that both techniques are similarly sensitive to an increase in surface roughness.

It is difficult to assess to what extent the deteriorating estimations are due to shadowing or to the reflectance function approximation. Section 7.3. is devoted to separating the effects of self shadows, cast shadows and surface roughness.

## 7.3 Robustness to shadowing

### 7.3.1 Motivation

Lambert’s law (and consequently the recovery algorithms) neglect both self and cast shadows as well as inter-reflections. In the computer graphics literature inter-reflections from diffuse surfaces are modelled using *radiosity* techniques. Even for simple scenes the computational cost is prohibitive. We are unaware of any convincing application to rough surfaces and do not model this effect. In this section we assess the effect shadowing has on the algorithms.

The degree of shadowing of a surface is a function of the roughness of the surface and of the zenith of the light source. We control the amount of shadowing by varying the surface roughness; the light zenith will be held constant. The effect of self and cast shadows is measured by comparing the algorithms’ ability to recover the surface for three cases:

- No shadowing.
- Self shadowing only.
- Self and cast shadowing.

This incremental approach allows us to isolate the effect of each phenomenon.

### 7.3.2 Method

We conducted an experiment where the surface recovery is repeated for increasingly rougher textures. The roughness of the three synthetic textures is varied in the range  $p_{rms} \in [0.1, 0.5]$  as in Section 7.2. The rendering process is divided into three stages: first the synthetic textures are rendered with Lambert’s law. We feed the proposed techniques with these shadow-free images and measure the recovery accuracy. Secondly, self shadows are modelled and added to the images. The recovery accuracy is again measured for the self-shadowed data. Finally, cast shadows are also added to the intensity images and the recovery performance is assessed.

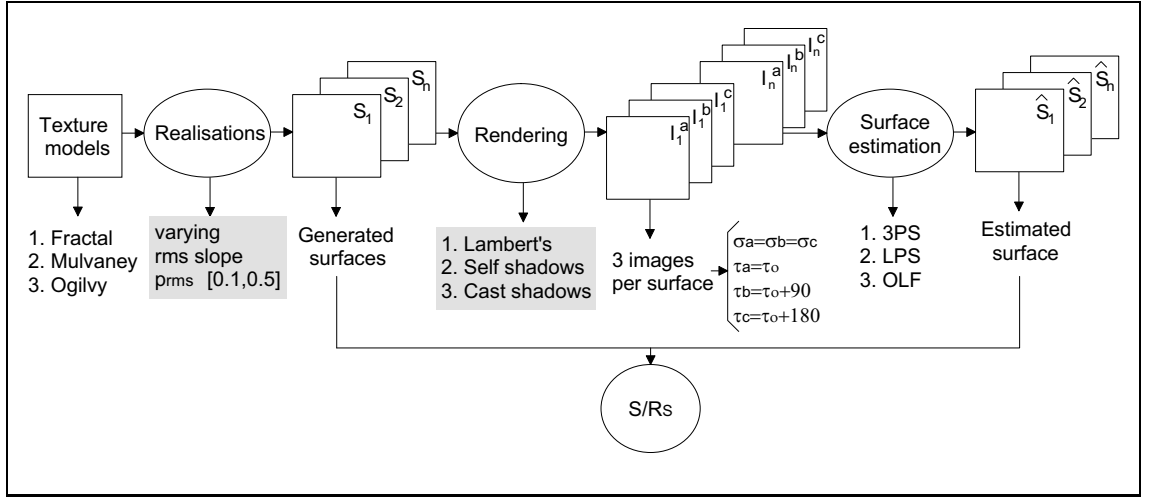


Figure 7.4: Procedure for assessing surface estimation against shadowing




Figure 7.5: 3-light PS performance against rms slope for synthetic Lambertian textures

To estimate the recovery accuracy we calculate the  $S/R_S$  ratio (equation 7.1). The  $S/R_S$  ratio is averaged across the fractal, Mulvaney and Ogilvy surfaces for each roughness. Figure 7.4 presents a scheme detailing the steps for this experiment.

### 7.3.3 Results and discussion

The performance of the 3-light PS is plotted in Figure 7.5 for the three cases: no shadows, only self shadows and both self and cast shadows.

The 3-light PS estimation is not dependent on roughness when shadows are not modelled. This technique assumes a perfect Lambertian surface and therefore it




Figure 7.6: Linear PS performance against rms slope for synthetic Lambertian textures

perfectly inverts the rendering process. Since shadows are neglected during the recovery, the more shadowing effects, the more obvious the deterioration. For the self-shadowed images, the  $S/R_s$  ratio drops about 3 dB relative to the no-shadow case. When considering both self and cast shadows, the  $S/R_s$  ratio drops by 5 dB.

Figure 7.6 shows the effects of surface roughness and shadows on the Linear PS estimation. Its performance drops by 0.5 dB when all shadowing effects are taken into account. This technique is quite robust to shadowing bearing in mind that the 3PS technique deteriorated by 5 dB in the same test. Although this method is more robust to shadowing, in absolute terms the 3PS performs more accurately over the roughness range (Figure 7.3).

The performance of the Optimal Linear Filter is plotted against roughness in Figure 7.7 for different levels of shadowing. When shadows are considered, the reflectance function approximated by the linear filter models shadows, i.e. the OLF -in a sense- can predict shadows. However, the true reflectance function becomes less linear and the estimation deteriorates. The filter's accuracy drops about 1.5 dB due to showing effects, i.e. it is relatively robust to shadows. Although for absolute values, the OLF gave in the worst estimate (Figure 7.3).

The experiments led us to the following conclusions. First, the 3-light PS algorithms is by far the most robust. Paradoxically, its invariance to surface roughness




Figure 7.7: Optimal Linear Filter (OLF) performance against rms slope for synthetic Lambertian textures

means that the effects of self and cast shadowing are most apparent with this algorithm. Secondly, although both the Linear PS and the OLF algorithms degrade significantly with increase roughness, shadowing is relatively a small factor in the degradation. Thirdly, the effect of cast shadowing is apparent only in 3-light PS and OLF.

## 7.4 Effect of surface reflectance

### 7.4.1 Motivation

The 3-light PS and the Linear PS assume a Lambertian surface. The aim of this assessment is to study the effect that surface specularities have on the recovery accuracy. Phong's model is used for rendering surfaces that depart from the Lambertian assumption. The diffuse term in Phong's model follows Lambert's law and can be approximately linear for relatively rough surfaces ( $p_{rms}, q_{rms} < 0.3$ ). The specular term can be thought of as a non-linear contribution added to the linear term. In this sense, Phong's model is also used to render surfaces whose reflectance departs from the linear function assumed by the OLF.

Two aspects of the specular reflection are taken into account. The first is the

amount of specularly reflected light relative to the total reflection. We secondly consider the nature of the specularity. Phong’s model produces different types of specular surfaces by modifying the shape and strength of the specular peak.

We intuitively expect a worsening of height estimation as the amount of specular reflection increases. On the other hand, predicting the estimators behaviour against specular nature is not easy. In Chapter 3, textures rendered with Phong’s model were compared to the same textures rendered with Lambert’s law (Section 3.8.2). As expected, Phong’s image departed from Lambert’s image when the specularity increased (Figure 3.9). We also analysed the differences between these intensity images when the specular peak shape was modified (Figure 3.10). For a wide low intensity specular peak, Phong’s intensity image is similar to Lambert’s image. As the specular peak is narrower and stronger, the similarities between Phong’s and Lambert’s images are drastically reduced. Finally, for very narrow specular peaks, the Phong’s image is again fairly similar to the Lambertian image. This is because most of the narrow specularities are reflected from the surface without reaching the viewer and do not affect the texture appearance. The illumination zenith is  $45^\circ$ , so the average mirror direction is also  $45^\circ$ . Since the viewer in the imaging system is straight above the surface, most of the narrow specularities are not seen. We expect the Linear PS and the 3-light PS response to specularities to be related to the similarity between Phong’s and Lambert’s images.

### 7.4.2 Method

In this section, we carry out two sets of experiments. Two parameters defined by Phong’s intensity expression (equation 3.10) are key factors for the assessment. These are the coefficient  $K_d$  that sets the percentage of diffusely reflected light relative to the total reflection, and the exponent  $n$  that models the specular peak shape.

With the first experiment we assess the estimation against the proportion of diffuse-to-specular contribution. For the rendering process, the coefficient  $K_d$  is

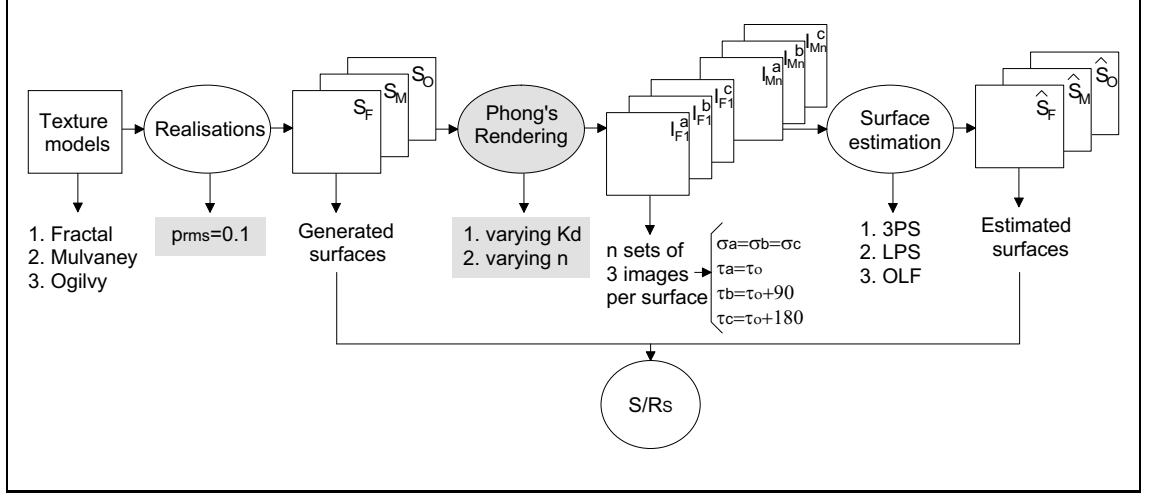


Figure 7.8: Procedure for assessing surface estimation against surface reflectance

gradually increased from 0 to 1 and the exponent  $n$  is fixed to 5.

In a second experiment, the shape of the specular peak is modified and its effect on surface estimation investigated.  $K_d$  is set to 0.9, i.e. the diffuse contribution is 90%. The specular nature is modified by varying  $n$  in the range  $[2, 26]$  in increments of 2. We model nearly diffuse surfaces by setting  $n$  to small values. Larger exponents give narrower specular peaks to model glossier surfaces.

In Chapter 3, we noted that varying  $K_d$  had a larger effect on surface reflectance than varying  $n$  (Section 3.8.2). Varying  $n$  mainly gives the impression of changing the light source size.

The roughness of the synthetic textures is set to  $p_{rms} = 0.1$  throughout the experiments. For a constant roughness and a fixed illuminant zenith, modelling shadows is not drastically affecting the techniques performance. Therefore, shadows are neglected in this case. A more detailed description of the experiments is shown in the scheme of Figure 7.8.

### 7.4.3 Results and discussion

First we assess the effect of increasing the surface specularities. As predicted, for an increasing diffuse contribution the estimation accuracy is improved (Figure 7.9). Although the improvement for the OLF is also noticeable, its sensitivity to specu-

larities is weaker.

For the considered range of  $K_d$ , the 3-light PS performance improves by more than 25 dB! Once the assumption of Lambertian reflectance is violated, the estimation is not accurate. In this particular case with  $n = 5$ , we cannot apply this technique to a surface whose diffuse reflection is smaller than 80% of the total. We cannot easily generalise this statement since the estimation depends on  $n$  as well as  $K_d$ . Nevertheless, Figure 7.9 still presents a very intuitive example of this technique's sensitivity to specularities. For an accurate recovery with the 3-light PS technique the surface has to be nearly Lambertian.

The Linear PS performance is similar to that of the 3-light PS. It is less sensitive to specularities, the  $S/R_s$  ratio increasing about 20 dB for the range of  $K_d$ . But its performance is still clearly unacceptable for relatively specular surfaces. The measured  $S/R_s$  ratio falls below the usual 10-dB threshold for  $K_d < 80\%$  and  $n = 5$ .

Although robustness against specularities depends on both  $K_d$  and  $n$ , we can conclude that in most cases the specular contribution should be less than 20%. For nearly Lambertian surfaces ( $K_d \approx 100\%$ ) the 3-light PS performance is best. In contrast, if the test surface has a wider range of reflectance characteristics ( $K_d \approx 80\%$ ), the Linear PS technique is recommended.

The OLF is comparatively robust to surface reflectance, improving by 4dB as  $K_d$  increased. This technique can cope with any amount of specularity since the estimation is accurate over the  $K_d$  range. With the OLF, we do not assume a Lambertian surface nor any other reflectance model, but the optimal linear function for the data. Fitting a linear function to Phong's model gives more accurate estimates than fitting Lambert's law.

Figure 7.10 shows the three estimators performance for a constant  $K_d$  and a varying  $n$ . The three estimates are accurate and over the 10dB threshold for a 90% of diffuse reflection. The specular image is close to the Lambertian image when  $n$  is small, thus the recovery is very accurate. As  $n$  increases, the surface is glossier and the techniques performance deteriorates. For large values of  $n$ , most of the






Figure 7.9: Recovery performance against percentage of diffuse reflection for synthetic textures

specularities are reflected by the surface without reaching the viewer. Thus the image rendered with Phong's model is again similar to the Lambertian image, and the estimations are improved. When the specularity takes the form of extremely narrow peaks, the Linear PS algorithm gives a better performance than the 3-light PS.

The Linear PS differs from the 3-light PS in the way they deal with non-linearities. By non-linearities we mean secondary and higher order terms of the Lambertian function as well as specularities. The 3-light PS algorithm estimates those non-linearities and compensates for them. Whilst the Linear PS algorithm neglects those non-linearities and assumes a linear reflectance. This latter approach proves to be more robust against specularities.

The  $S/R_s$  ratio for the Optimal Linear Filter increases by about 3dB as the exponent  $n$  increases. The improvement shown by both the Linear PS and the OLF for a large  $n$  may be related to the validity of a linear reflectance function.




Figure 7.10: Recovery performance against specular peak shape for synthetic textures

## 7.5 Robustness against noise

### 7.5.1 Motivation

Images often contain a certain level of random noise generated either during the imaging process or due to a noisy communication link. It is important that any surface recovery technique can operate successfully with a given level of noise. To test the noise performance of the proposed schemes, various levels of noise were added to the textured images. Noise is considered to be additive white and Gaussian.

McGunnigle [59] looked for a justification of the form and the amplitude of the noise model. He found support either in the literature or through experiments and simulation. He argued that the temporal noise associated with the subsampled image was reasonably approximated as white noise of variance 1.17. He estimated this variance for a certain dataset, i.e. his noise model was not general. Temporal noise was due to mechanisms such as dark noise, shot noise or jitter effects. McGunnigle also found a justification to model quantisation noise as white noise. To make our case more general, no specific level of noise is searched for. So its amplitude is varied throughout the simulations and the maximum acceptable signal-to-noise ratio is measured.




Figure 7.11: Fractal (left), Mulvaney (centre) and Ogilvy (right). Top row: no noise. Second row:  $S/N = 10dB$ . Bottom row:  $S/N = 0dB$ .

### 7.5.2 Method

The experimental data are the synthetic surfaces, several realisations of a white Gaussian process and a Lambertian rendering process. The surface roughness is fixed to  $p_{rms} = 0.1$ . We render the textures to get noise-free images. Noise is then added as a white Gaussian process with zero mean and a variance dependent on the required signal-to-noise ratio. Six levels of noise were used, ranging from a barely visible 25 dB signal-to-noise ratio to a very obvious 0 dB.

Figure 7.11 shows the images of the three textures with different levels of added noise. The top row images are free of noise; on the second row the  $S/N$  is 10dB; and on the bottom row, the  $S/N$  is 0 dB. The images were rendered under the same illumination conditions.

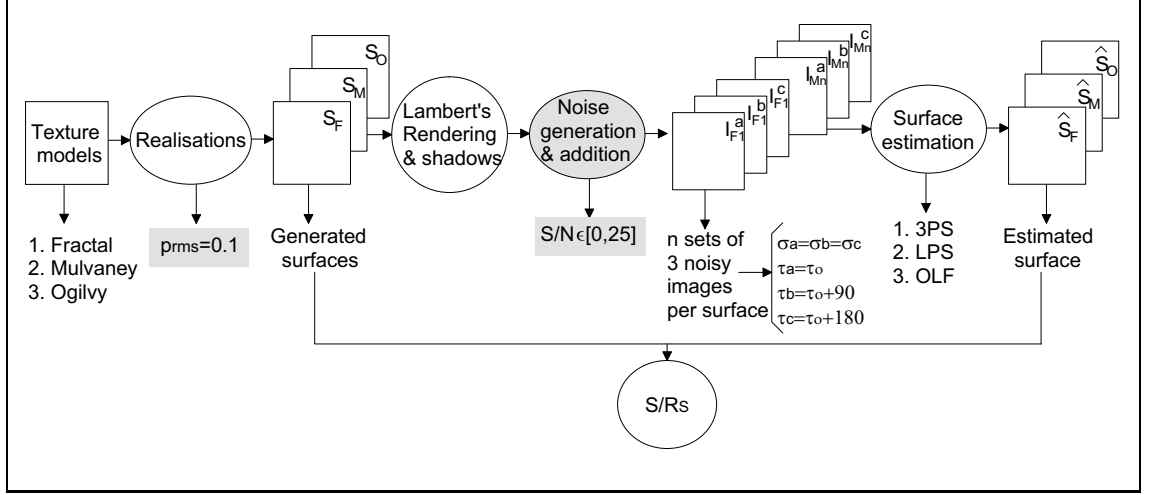


Figure 7.12: Procedure for assessing surface estimation against noise

The experimental procedure is shown in Figure 7.12

### 7.5.3 Results and discussion

Figure 7.13 shows the performance of our height recovery techniques against noise. The recovery is almost unaffected for a wide range of  $SNR$  of the images. Only for very noisy images, with  $SNR$  smaller than 15 dB, is a deterioration in performance really noticeable. The algorithms show similar sensitivity to noise since their response follows the same curve. The total  $S/R_s$  ratio fall is less than 9 dB in the whole  $SNR$  range for all methods.

Although all techniques are analogously robust to noise, the 3-light PS performance is superior. The measured  $S/R_s$  ratio for the 3-light PS and the Linear PS is above the 10 dB threshold for the whole range of  $SNR$ . For the OLF, the  $S/R_s$  ratio is below the threshold for images with the minimum  $SNR$ . Thus this technique can only cope with images whose  $SNR$  is above 5 dB.

## 7.6 Discussion

In Chapter 5 we identified potential shortcomings in the schemes due to their assumptions about the surface topography and reflectance (Section 5.4). In this chap-




Figure 7.13: Recovery performance for synthetic textures against noise

ter we measured the effect of violating these assumptions. We applied our algorithms to textured surfaces with a wide range of different characteristics. Simulation was necessary since the test surface features had to be closely controlled. For purposes of comparison, the Benchmark 3-light PS technique was also considered in the assessment. The effect of surface roughness and reflectance and the robustness against shadowing and noise were investigated.

The benchmark 3-light PS was the most robust to surface *roughness*, giving a signal-to-residue ratio greater than 20dB for  $p_{rms}, q_{rms} \leq 0.35$ . The Linear PS and the OLF are much less robust with SRR greater than 10dB over a similar range. This covers a wide range of textures, thus we consider the novel techniques suitable for rough surface recovery.

The Optimal Linear Filter was found the most robust against *specular reflection*. The Linear PS and the 3-light PS were very sensitive to specularities. Whenever ground truth data is not available, the 3-light PS technique should be applied to purely Lambertian surfaces. The Linear PS technique is recommended for surfaces that are not perfectly diffuse and show specular behaviour.

The 3-light PS algorithm was found to be most robust to *shadowing*. The Linear PS and OLF algorithms gave similar results, although the Linear PS was slightly better. The OLF is not good at predicting shadows as expected (see Section 5.3.3).

All the algorithms were able to perform well under very severe levels of *noise*,

the 3-light PS being best.

We should bear in mind some of the limitations that the simulation presents. These limitations extend to the surface models, the surface-to-image transfer function and the model for noise.

- For our surface models every pixel corresponds to a surface facet, whilst in reality it corresponds to a small area of the continuous surface function. Moreover we do not know how representative these models are of real textures.
- The reflectance function is too simple to be realistic. The surface reflectance is assumed to follow either Lambert's law or Phong's model. Shadows are modelled as complete absence of intensity and inter-reflections are neglected.
- Noise is assumed to follow a random process which is additive, white, Gaussian and independent of the signal.

Furthermore this assessment is insufficient since we investigated the techniques' shortcomings individually but some effects are correlated. For real data, these aspects are observed simultaneously. In Chapter 8 and Chapter 9, the assessment of our surface estimators is extended to real textures.

## Chapter 8

# Assessment of Surface Estimation: Experiment

---

Testing the proposed recovery techniques on real data is essential. With simulation we made assumptions about the surface models, the noise model and the surface-to-image transfer function. Assessing our techniques on real data removes these assumptions. Furthermore, a larger number of surface types are taken into consideration in this chapter.

Because ground truth data is not available, the assessment of our techniques becomes difficult. The Optimal Linear Filter cannot be assessed on our real textures because it must be trained on ground truth data. When the OLF is trained on our synthetic models, the filter weights are equal to those of the LPS method (see Section 5.2.1). Therefore, only the 3-light Photometric Stereo and the Linear Photometric Stereo are evaluated in this chapter.

One way to assess the accuracy of recovery is through rendering; the recovered surface is synthetically rendered, and the resulting image is compared to the original surface image. The performance of our methods is tested on seventeen real textures. We assess the scope of the techniques' shortcomings — identified in Chapter 5. We study the robustness of the recovery to surface roughness and surface discontinuities for several surface types and reflectance functions. Four exceptional textures that violate our assumptions about the surface are also considered.

We study the possibility of using our estimators for image prediction purposes. For this second experiment the required accuracy is not as high as for surface recovery. Furthermore, a wider range of lighting conditions is considered. The accuracy of image prediction to changes in the light azimuth and zenith is investigated.

## 8.1 Assessment of the recovery through rendering

This chapter starts by describing the method used to assess the surface recovery through rendering. The experiment, the data set, the accuracy metric and the weaknesses of the assessment are detailed. Next the results are presented and discussed.

### 8.1.1 Experimental method

#### Description

We carry out an experiment where the textured surface is recovered from its images using the proposed algorithms. To assess the recovery we compare the rendered surface estimate to original intensity images. Ideally, we should render the shape estimate using the surface reflectance properties. Instead, we assume that the reflectance of the test surfaces is perfectly modelled by Lambert's law. Therefore, any difference between the synthetic and the real image is assumed to be due to an inaccurate surface estimate. This assumption is in fact a deficiency of our assessment procedure. However, estimating the surface reflectance properties is out of the scope of this thesis.

The surface heightmap is estimated for images lit with the smallest zenith, i.e.  $45^\circ$  or  $60^\circ$  depending on the sampling of the texture, and azimuth angles in  $90^\circ$  increments (see Section 2.4.3).

The estimate is rendered for the same lighting conditions as the real images. For a more robust assessment, we consider several different lighting conditions. We evaluate the recovery for two light zenith angles and azimuths at  $30^\circ$  increments, i.e. twenty-four images. We do not assess the recovery on the images that were used for surface estimation. We are therefore evaluating the algorithms' ability to *predict* the surface appearance.

The accuracy of prediction is measured using the signal-to-residue ratio (SRR):



$$\frac{S}{R_i} = 10 \log \left[ \frac{\text{var}[i(x, y)]}{\text{var}[i(x, y) - \hat{i}(x, y)]} \right] \quad (8.1)$$

where  $i(x, y)$  is the real image and  $\hat{i}(x, y)$  is the rendered heightmap. The SRR is estimated for each lighting condition. It is then averaged to give a single measurement of accuracy:

$$\frac{S}{R_i} = \frac{1}{n} \sum_{j=1}^n 10 \log \left[ \frac{\text{var}[i_j(x, y)]}{\text{var}[i_j(x, y) - \hat{i}_j(x, y)]} \right] \quad (8.2)$$

where  $i_j$  is the  $j^{\text{th}}$  real image,  $\hat{i}_j$  is the  $j^{\text{th}}$  rendered heightmap and  $n$  is the number of test images considered for the assessment. For ease of discussion, a surface is considered to be accurately recovered if the SRR of equation 8.2 is over 10dB.

## Data

We use the test surfaces described in Section 2.4. The real textures have several surface height and reflectance functions and differ in their directionality and roughness.

The *fracture* textures are a collection of Lambertian fractal-like surfaces of increasing roughness (Figure 2.5 (a), (b), (c), (d) and (e)). This set allows us to study the algorithms' robustness to surface roughness.

Similarly, we can investigate the effect of violating the Gaussian assumption using the *deposit* surfaces (Figure 2.5 (f), (g), (h), (i), (j) and (k)).

The *sand-ripple* textures have similar topographies but the amplitude and frequency of the waveforms vary (Figure 2.6 (a), (b) and (c)). This translates into a variation of roughness, but it also means that the surface pattern varies from surface to surface. Thus we can study the effect that modifying the surface topography has on the estimation.

The *anaglypta* textures violate the Lambertian assumption and present several topographic functions (Figure 2.6 (e), (f) and (g)).

Four other textures were included in the database because they severely violate the surface assumptions. The *irregular stripes anaglypta* is a non-Lambertian rough surface (Figure 2.6 (d)). The *repetitive primitive* is phase-rich, and its image is heavily shadowed (Figure 2.6 (i)). The *textile* is a discontinuous surface and does not have constant albedo (Figure 2.5 (l)). The *sand paper* shows a very specular reflection for some surface facets (Figure 2.6 (h)).

## Rendering shadows

Our model to render shadows is very simple. We only consider self-shadows, i.e. no cast shadows are modelled. Furthermore, we assume no ambient illumination or inter-reflections, so that the shadowed facets have intensity zero. This simple model is not realistic. For shallow illumination, the amount of shadowing increases and the rendering inaccuracies become more significant. To minimise the error introduced during rendering, we do not consider shallow lighting conditions. Only non-Lambertian smooth surfaces, such as the anaglyptas, are assessed for the largest zenith ( $75^\circ$ ).

## Constant albedo

Although our surfaces are assumed to have constant albedo, they are likely to present a small deviation in the albedo map. We do not assume a constant albedo for rendering purposes. Instead, the albedo map is estimated using the 3-light Photometric Stereo algorithm and accounted for. There are two reasons to incorporate the albedo information of the real surfaces in the assessment procedure:

- To improve the validity of our assessment we should use the most realistic rendering process. We should consider all the reflectance information that we can estimate from the test surface, and its albedo is available to us.
- The 3-light Photometric Stereo algorithm estimates albedo on a pixel by pixel basis. In contrast, the Linear Photometric Stereo does not allow for variations

in albedo. The albedo variation, if exists, will be incorporated into the surface estimate. By rendering with constant albedo, we would be favouring the Linear Photometric Stereo because the surface estimate accounts for albedo variations. The error due to the rendering process would be larger for the 3-light PS estimate than for the Linear estimate.

However, the estimated albedo is a possible source of errors. The estimation errors and noise contained in the albedo map are added to the images during rendering. This means that errors due to the 3-light Photometric Stereo estimation affect the rendered Linear PS estimate.

The accuracy of prediction when using albedo information to render the Linear estimate is improved by 0.14dB on average (see Appendix B, Section B.1). Thus we chose to use the albedo map for rendering both the Linear and the 3-light PS estimates.

It is worth noting that the albedo is only used for rendering during the assessment procedure. It is not used for improving the estimation in any way.

### **Incident intensity**

When rendering the estimated surface, we simulate the lighting conditions of the real image. We know the light source position but not the incident intensity value. This value,  $I_0$ , is a constant multiplicative factor in Lambert's expression (equation 3.3). To calculate the incident intensity, we render the surface heightmap for  $I_0 = 1$ . We remove the mean of the synthetic and the real image, and calculate  $I_0$  as:

$$I_0 = \sqrt{\frac{\text{var}[i(x, y)]}{\text{var}[\hat{i}(x, y)]}} \quad (8.3)$$

The incident intensity for a certain light zenith is the average of the  $I_0$  calculated for the available test images lit with that zenith.

Texture	pixel-width [mm]
<i>Fractures, deposits and sand ripples</i>	0.1
<i>Anaglyptas</i>	0.2
<i>Irregular stripes anaglypta</i>	0.15
<i>Repetitive primitives</i>	0.15
<i>Sand paper and textile</i>	0.1

Table 8.1: Pixel-width of the image of the real textures [mm]

## Height units

The surface estimate is given to us in pixel-width. We estimate the dimensions of a pixel by measuring the coordinates of two points of known separation on a surface that lies on a plane parallel to the image plane. Table 8.1 shows the pixel-width for the textures considered.

In Section 4.5.2, we argued that the absolute distance from the viewer to the surface cannot be recovered. This is the case with both the Linear PS and the 3-light PS techniques.

### 8.1.2 Experimental results

The recovery techniques are tested on seventeen real textures which meet the surface assumptions to varying degrees. Therefore we can separately study the robustness of our estimates to surface discontinuities, roughness, topography and reflectance. Four special textures that severely violate the surface assumptions are also considered. This gives a total of twenty-one real textures. Accuracy of *recovery* is expressed in terms of the accuracy of *prediction* of the surface appearance.

A more detailed description of the assessment results is included in Section B.2 of Appendix B.




Figure 8.1: Recovery accuracy of the Linear PS and the 3-light PS techniques

### Assessment of surface estimation through rendering

Figure 8.1 shows the accuracy of recovery for the real textures using both the Linear estimate and the 3-light PS estimate. The name that the textures are given in the graph corresponds to the figure where they are plotted in Chapter 2. For instance, the texture *5b* is the *moderate roughness fracture* of Figure 2.5 (b).

The performance of the Linear PS technique is inferior to that of the 3-light PS algorithm for most textures. Only four out of seventeen textures were accurately recovered with the Linear method. In contrast, the 3-light PS technique accurately estimates ten of the considered textures.

We argued that the rendering assessment procedure is limited (the surface reflectance is assumed to be Lambertian) and it introduces errors (errors in the albedo map estimate). Therefore, we can infer that the *measured* accuracy of prediction is a lower bound on the *actual* accuracy of recovery.

We conclude that the 3-light PS technique is successful in recovering most test textures and that the Linear estimate has to be improved.

As an example of an accurate estimation, Figure 8.2 (a) shows a real image, (b) the rendered 3-light estimate and (c) the rendered Linear estimate. The texture is




Figure 8.2: *Moderate roughness fracture* imaged for  $\sigma = 45^\circ$  and  $\tau = 300^\circ$

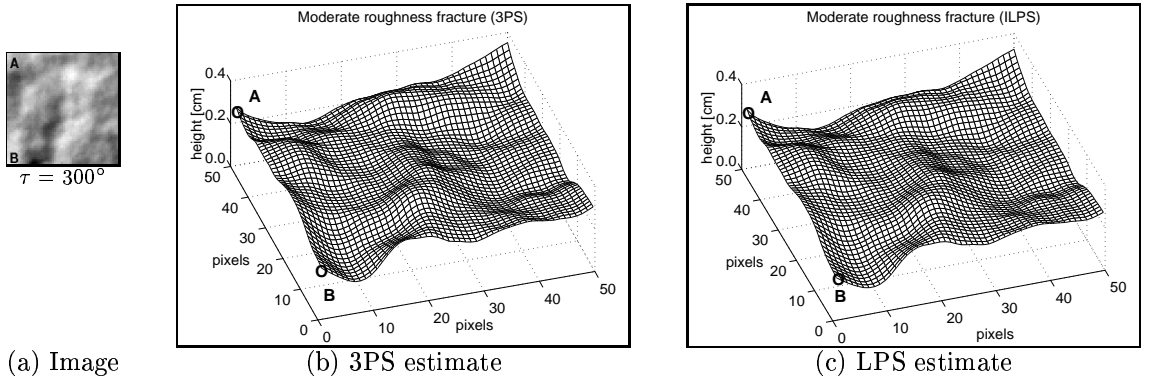


Figure 8.3: Estimated *moderate roughness fracture* surface

the *moderate roughness fracture* imaged with a light zenith of  $45^\circ$  and an azimuth of  $300^\circ$ . The accuracy of prediction for this particular image of the texture is 12.51dB for the 3-light PS technique and 9.97dB for the Linear technique. We can get a feeling for what 10dB mean by visually comparing the rendered Linear estimate to the real image: the test image is well approximated by the synthetic prediction.

The estimated heightmap for the *moderate roughness fracture* is plotted in Figure 8.3. We selected an area of 50 by 50 pixels from the top left corner of the surface image (Figure 8.3 (a)). The surface heightmaps recovered with both techniques are plotted in Figure 8.3 (b) and (c). Both estimates are similar to each other and seem to correctly follow our idea of the surface shape.

### Robustness to surface roughness

To assess the robustness of our techniques to roughness we study the recovery performance for the *fractures*, which cover a wide range of roughnesses. Table 8.2 shows

Texture	Figure	LPS	3PS
<i>Rough fracture</i>	<i>2.5 (a)</i>	7.22	8.52
<i>Moderate roughness fracture</i>	<i>2.5 (b)</i>	9.69	13.13
<i>Gentle roughness fracture</i>	<i>2.5 (c)</i>	8.00	12.71
<i>Smooth cured fracture</i>	<i>2.5 (d)</i>	7.60	13.37
<i>Smooth patterned fracture</i>	<i>2.5 (e)</i>	6.41	12.84

Table 8.2: Average accuracy of recovery for fractures [dB]

the algorithms’ ability to predict the appearance of the *fracture* estimates. We expect the photometric estimates to deteriorate for rougher surfaces since shadowing effects increase. Also non-linear effects become dominant further degrading the Linear PS estimation. For the *rough fracture* the prediction accuracy is poor, below the 10dB threshold. For smoother surfaces, the rendered 3-light PS estimate is closer to the test images. The surface is accurately estimated with the 3-light PS. In contrast, the SRR is never over 10dB for the Linear method. Even for the smoother surfaces, the Linear PS fails to accurately predict the test images.

The *rough fracture* of Figure 2.5 (a) is chosen as an example of an inaccurate estimate. Figure 8.4 shows a test image of the fracture (a), the rendered 3-light estimate (b) and the rendered Linear estimate (c). The estimated albedo map mainly captures shadows and highlights for a constant albedo surface like this (Figure 8.4 (d)). The image prediction accuracy is 7.03dB with the 3-light PS technique and 6.84dB with the Linear technique. The difference between the real image and the rendered estimates is obvious. However, the residue images — calculated by subtracting the rendered heightmap from the real image — are shown in Figure 8.4 (e) and (f). Shadows and highlights are not properly predicted, although they are encoded to a certain extent in the albedo. Furthermore, the rendered estimates fail to show the degree of roughness of the texture.

Figure 8.5 (a) shows the 50x50 pixels at the top left corner of the surface image. Figure 8.5 (b) and (c) are the estimated heightmaps. These two estimates are




Figure 8.4: *Rough fracture* imaged for  $\sigma = 45^\circ$  and  $\tau = 300^\circ$

fairly consistent with each other. Although, considering the poor quality of recovery measured, they do not seem to accurately model the surface.

### Robustness to surface discontinuities

We assume that the test surfaces follow a Gaussian distribution of heights. How is the recovery affected by a violation of the Gaussian assumption? We test the novel algorithms on six *deposit* surfaces (Table 8.3). The lighter the deposit, the less Gaussian the surface height distribution [33]. Both estimates deteriorate for more

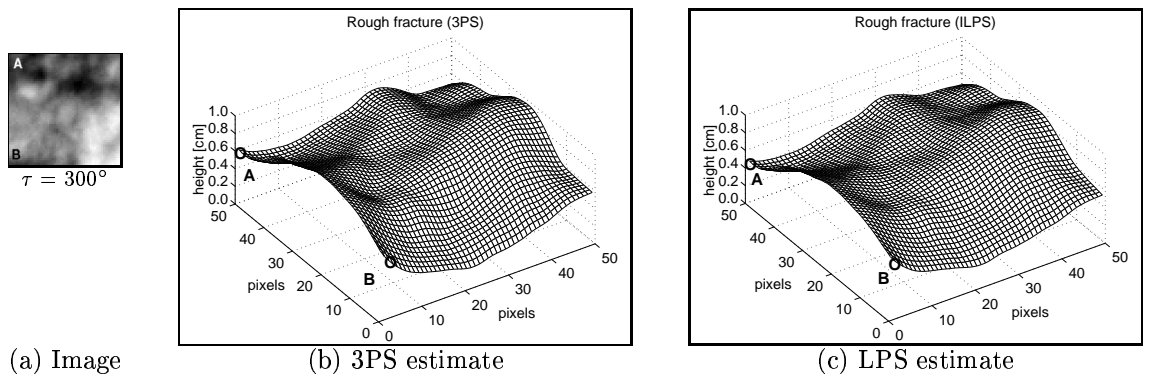


Figure 8.5: Estimated *rough fracture* surface



Texture	Figure	LPS	3PS
<i>Moderate rough deposited fracture</i>	<i>2.5 (f)</i>	8.79	10.16
<i>Rough deposited fracture</i>	<i>2.5 (g)</i>	8.86	10.36
<i>Heavy deposited surface</i>	<i>2.5 (h)</i>	8.86	9.92
<i>Medium deposited surface</i>	<i>2.5 (i)</i>	8.36	9.14
<i>Light deposited surface</i>	<i>2.5 (j)</i>	8.28	8.67
<i>Sparse deposited surface</i>	<i>2.5 (k)</i>	7.04	7.04

Table 8.3: Average accuracy of recovery for deposited surfaces [dB]

sparsely deposited surfaces. The 3-light PS estimate is accurate for rough deposited textures and heavily deposited surfaces. The prediction accuracy for the Linear PS estimate is less than 10dB for all surfaces. However, the 3-light PS estimate is more sensitive to surface discontinuities than the Linear estimate. The SRR decreases by 3.32dB from the most accurate estimate (*rough deposited fracture*) to the least accurate estimate (*sparse deposited surface*). This deterioration is only 1.82dB for the Linear technique.

Figure 8.6 (b) and (c) show the recovered heightmaps of the *heavily deposited texture*. The plotted heightmaps correspond to an area of 50x50 pixels at the top left corner of the surface image (Figure 8.6 (a)). The estimates are similar to each other and seem to properly reproduce the low frequency undulations of the surface. However, the high frequency height variation due to most recent deposits are not rendered.

The 3-light PS technique performs better than the Linear method for these textures. However, it is more sensitive to surface discontinuities.

### Surface estimation for several topographic and reflectance functions

To properly test the robustness of our methods to surface topography and reflectance, we would ideally modify the surface topography in a controlled manner whilst keeping the same reflectance and vice versa. This is possible in simulation

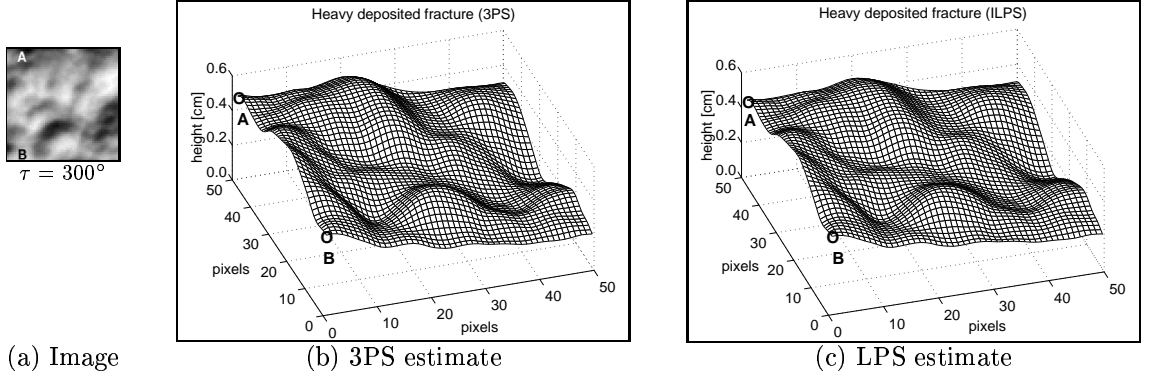


Figure 8.6: Estimated *heavy deposited texture*

but not with real data. No set of textures that gradually differ in their topography or reflectance was available. Thus, we recover several surfaces with different topographic and reflectance functions. We cannot generalise the results obtained for these textures to other textures. However, they give an intuitive idea of the algorithms' ability to recover surfaces with similar features to the considered test surfaces.

Three *sand ripples* and three *anaglypta* textures are considered. Table 8.4 shows the recovery performance for both estimators. The accuracy of prediction when rendering the Linear PS estimate is very close to 10dB for all surfaces. In contrast, the 3-light PS estimate is accurate for the Lambertian sand ripples but not for the non-Lambertian anaglypta surfaces. The 3-light PS estimate is less robust to specularities than the Linear estimate. We reached the same conclusion when we assessed our recovery methods through simulation (Chapter 7).

We should note that the directional surfaces have their grain aligned to the  $90^\circ$  direction. One of the estimation images was lit with an azimuth of  $90^\circ$ , so non-linear terms are dominant for this image. We would expect the Linear estimate to deteriorate under these conditions. Table 8.4 shows that the Linear estimation is poorer than the 3-light estimation for the directional sand ripples. The accuracy of recovery for directional surfaces is further investigated in Section 8.2.2.

Texture	Figure	LPS	3PS
<i>Sand-rippled surface</i>	2.6 (a)	9.53	11.75
<i>Sand-rippled surface (higher frequency)</i>	2.6 (b)	11.01	11.33
<i>Smoothed sand-rippled surface</i>	2.6 (c)	10.11	10.59
<i>Net anaglypta</i>	2.6 (e)	10.04	7.68
<i>Rippled anaglypta</i>	2.6 (f)	10.34	8.82
<i>Rice grain shaped anaglypta</i>	2.6 (g)	9.68	9.09

Table 8.4: Average accuracy of recovery for sand-rippled and anaglypta surfaces [dB]

### Surface estimation for special textures

In this section we assess our algorithms on textures that severely violate the surface assumptions. The *irregular stripes anaglypta* has a hybrid reflectance and the *repetitive primitives surface* has many discontinuities. Both are very rough, thus their images show significant shadowing. Our recovery techniques give poor results in predicting the image appearance for these textures (Table 8.5). It is worth noting that shadowing affects the surface estimation as well as the rendering process. The proposed techniques are not suitable for estimating such rough surfaces.

The *textile* surface is a discontinuous rough texture with non-constant albedo. The accuracy of prediction is poor for both recovery algorithms (Table 8.5). The Linear method assumes constant albedo, thus this estimate is poorer than the 3-light estimate. For most applications, these results are unacceptable.

The last special case considered is a *sand paper* surface which has specular reflectance. The accuracy of prediction of the texture appearance is very low (Table 8.5). The estimation performance deteriorates for specular surfaces since a Lambertian reflectance is assumed. Furthermore, the Lambertian rendering that is applied to the surface estimate is an inappropriate model of the real reflectance.

The *repetitive primitive texture* is used to show the estimation failure (Figure 8.7). For the test image with zenith  $45^\circ$  and azimuth  $330^\circ$  the accuracy of image

Texture	Figure	LPS	3PS
<i>Irregular stripes anaglypta</i>	2.6 (d)	5.13	1.92
<i>Repetitive primitive surface</i>	2.6 (i)	3.84	5.76
<i>Textile</i>	2.5 (l)	2.89	4.43
<i>Sand paper</i>	2.6 (h)	0.57	-1.72

Table 8.5: Average accuracy of recovery for special cases [dB]




Figure 8.7: *Repetitive primitive texture* imaged for  $\sigma = 45^\circ$  and  $\tau = 330^\circ$

prediction of the 3-light PS technique is 7.52dB. Similarly, the Linear estimate does not accurately model the surface appearance (3.59dB). One of the main sources of error is due to the the absence of shadows in the synthetic images. Furthermore, the rendered estimates seem to be blurred in comparison to the real image. Apart from the poor SRR measured for these images, they still resemble the test image. This means that for the only purpose of image prediction, the SRR threshold could be relaxed to a lower value than 10dB.

We assess the surface estimates by studying the heightmaps obtained with both techniques (Figure 8.8). Only a fraction of the test image is recovered: an area of 50x50 pixels belonging to the top left corner. The surface estimates are not as similar to each other as in previous examples. The heightmaps correspond to a continuous surface, whilst the real surface is not continuous. In Chapter 7 we concluded that the 3-light PS technique was more robust to surface roughness than the Linear method. This is also observed when comparing the estimates; the 3-light PS heightmap seems to follow the barley grains shape more accurately than the Linear estimate.

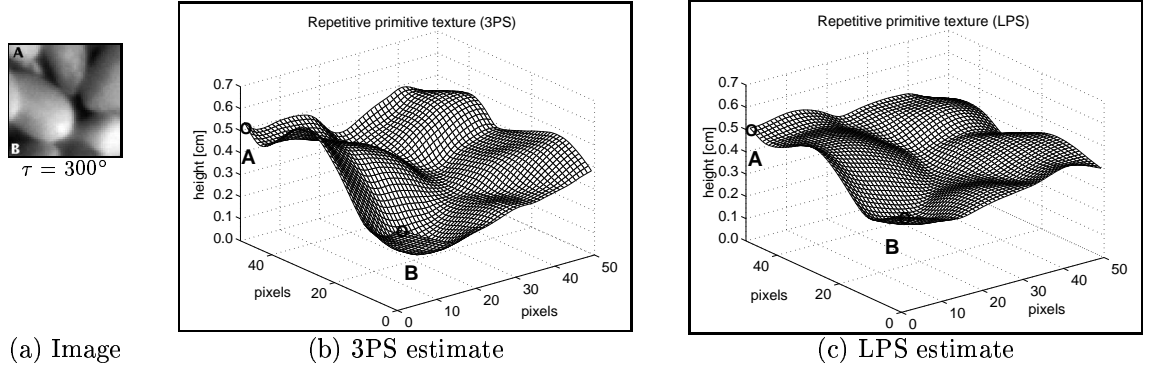


Figure 8.8: Estimated *repetitive primitive texture*

## 8.2 Surface estimation for image prediction purposes

In many application the only interest is the surface appearance and not the surface shape. A supervised classifier and an estimator of the light position are examples. The input data to the classifier is the surface image taken with unknown lighting conditions. Thus, the surface estimate can be used to predict the surface appearance for training the classifier or to estimate the source position.

This assessment is different to the one proposed in Section 8.1 since image prediction is the ultimate aim and not an assessment tool. In Section 8.1.2 we argued that for the only purpose of image prediction the SRR could be relaxed from the threshold of 10dB. We now assess the accuracy of image prediction for a wider range of illumination conditions. Importantly, the robustness of our prediction to changes in the lighting conditions is investigated.

### 8.2.1 Experimental method

#### Description

The surface recovered with our photometric methods is rendered and compared to the test image. We render the surface estimate using Lambert's law for the same lighting conditions as the real image. The surface is estimated for images with the smallest considered zenith and azimuth angles in  $90^\circ$  increments.

We study the robustness of image prediction when varying the azimuth angle and for three zenith angles. The reflectance function of a directional Lambertian surface could be dominated by non-linear terms depending on the light azimuth position. Thus we investigate separately the accuracy of image prediction in directional and isotropic surfaces.

For rendering the recovered surface shapes, we use the albedo map estimated with the 3-light PS technique. We therefore account for the small variation of the surface albedo. We also use the estimated intensity value for rendering (equation 8.3).

We measure accuracy of image prediction as the SRR for each available image of the test texture (equation 8.1). The SRR threshold is now set to 6dB.

## Data

We use the data set described in Chapter 2. However, we exclude the four textures that we categorised as being special cases. These textures were included in the database to test the limits of surface recovery. The recovery techniques showed poor performance for these surfaces. Therefore, they are not considered for assessing the accuracy of prediction.

## Rendering shadows

When assessing the accuracy of recovery, we avoided those lighting conditions for which the rendering process was not realistic. In this way, shallow illuminations were not considered because shadows were not properly modelled. In contrast, when assessing the prediction of the surface appearance we are somehow assessing the realism of the rendering process. Although the model to render shadows is very basic, we should take into account shallow lighting conditions.

The three sampling zenith angles,  $45^\circ$ ,  $60^\circ$  and  $75^\circ$ , are considered in this investigation (see Section 2.4.3). This gives a total of thirty-six images per texture. We should note that the anaglypta surfaces are imaged with only two zenith angles ( $60^\circ$  and  $75^\circ$ ).

<i>Technique</i>	<i>LPS</i>			<i>3PS</i>		
<i>Test zenith</i>	$\sigma = 45^\circ$	$\sigma = 60^\circ$	$\sigma = 75^\circ$	$\sigma = 45^\circ$	$\sigma = 60^\circ$	$\sigma = 75^\circ$
<i>Fractures (5)</i>	8.76	8.28	8.06	14.93	13.02	7.23
<i>Deposits (6)</i>	9.94	8.44	8.36	11.91	9.50	3.18

Table 8.6: Accuracy of the prediction of the surface appearance for isotropic surfaces [dB]

## 8.2.2 Results and discussion

### Effect of varying the light zenith for isotropic surfaces

Table 8.6 shows the averaged accuracy of image prediction for the *fracture* and the *deposit* surfaces. We consider separately the prediction accuracy for three rendering zenith angles. A more detailed table with the SRR per texture per zenith angle is included in Appendix B (Section B.2).

The rendered heightmap is closest to the test images for the estimation zenith ( $45^\circ$ ). As the rendering zenith departs from the estimation zenith, the test images differ from the estimation images. Any inaccuracy in the estimated heightmap is more noticeable as the rendering conditions differ from the estimation conditions.

We should also take shadowing effects into account. For shallow illumination, shadowing increases and the prediction deteriorates. For the 3-light PS technique, the prediction accuracy drops an average of 8 dB as the zenith increases. Why are shadowing effects not so apparent in the predictions of the Linear technique? In Section 5.3.2, we argued that the linear term in the Lambertian intensity expression was dominant for shallower zenith angles. We expect the rendered Linear estimate to be closer to the near linear test image observed for raking light. The effect of having a more linear surface image evens out with shadowing effects, and the prediction accuracy is not so dependent on the light zenith. Although the Linear PS seems more robust to shadowing effects, this does not necessarily mean a more accurate surface estimate.

As an example we compare the images of the *heavily deposited surface* to its




Figure 8.9: *Heavy deposited surface*: Real image (left); 3PS estimate (centre); LPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 150^\circ)$ ; bottom row:  $(\sigma, \tau) = (75^\circ, 150^\circ)$

estimate rendered with zenith  $45^\circ$  and  $75^\circ$ . Figure 8.9 (a) and (d) are the two test images. Figure 8.9 (b) and (e) are the two rendered 3-light PS estimates. And Figure 8.9 (c) and (f) are the rendered Linear estimates. All the images have the same light azimuth ( $150^\circ$ ). Both estimates rendered with zenith  $45^\circ$  are very close to the corresponding test image. In contrast, the Linear PS estimate rendered for  $\sigma = 75^\circ$  is closer to the test image than the 3-light PS estimate. When the illumination is shallow, the image is predicted more accurately from the Linear estimate than from the 3-light PS estimate.

The prediction of the appearance of the *rough deposited fracture* is investigated for a range of light azimuths and three zeniths. Figure 8.10 shows the performance for the 3-light PS estimate. The SRR of the three estimation images is over 22dB. However the prediction deteriorates for an azimuth angle different to the estimation angles — for an azimuth of  $270^\circ$ , the accuracy drops to 10dB. Similarly, as the rendering zenith departs from the estimation zenith, the prediction deteriorates drastically.






Figure 8.10: *Rough deposited fracture*: Accuracy of prediction when rendering the 3-light Photometric Stereo estimate

Figure 8.11 shows the accuracy of prediction for the Linear technique. The estimation images are predicted with a SRR of 12dB. The accuracy is highest for the estimation zenith curve ( $\sigma = 45^\circ$ ). The dependency on the zenith angle is not so strong for this estimation.

### Effect of varying the light zenith for directional surfaces

Directional surfaces are especially difficult to recover. When the light azimuth direction is parallel to the surface grain, non-linear order effects are dominant. Under these conditions the rendered Linear Photometric Stereo estimate may not approximate the test image well. Table 8.7 shows the averaged accuracy of prediction for the *sand-ripples* and *anaglypta* surfaces<sup>1</sup>. Table B.1 in Appendix B includes a more detailed description of the assessment results.

The performance of the 3-light PS technique does not deteriorate for shallow illumination. To better explain this effect we analyse the case of the *sand-rippled surface* (Figure 2.6(b)). Figure 8.12 shows the real images and the rendered surface estimates for two zenith angles:  $45^\circ$  and  $75^\circ$ . The main difference between the

<sup>1</sup>Note that the estimation zenith for the *sand ripples* is  $45^\circ$  and for the *anaglyptas* is  $60^\circ$ .




Figure 8.11: *Rough deposited fracture*: Accuracy of prediction when rendering the Linear Photometric Stereo estimate

<i>Technique</i>	<i>LPS</i>			<i>3PS</i>		
<i>Test zenith</i>	$\sigma = 45^\circ$	$\sigma = 60^\circ$	$\sigma = 75^\circ$	$\sigma = 45^\circ$	$\sigma = 60^\circ$	$\sigma = 75^\circ$
<i>Sand-ripples (3)</i>	10.93	11.33	11.24	12.36	13.17	9.97
<i>Anaglyptas (3)</i>	—	11.84	10.17	—	9.20	10.16

Table 8.7: Accuracy of the prediction of the surface appearance for directional surfaces [dB]




Figure 8.12: *Sand-rippled surface*: Real image (left); 3PS estimate (centre); LPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 150^\circ)$ ; bottom row:  $(\sigma, \tau) = (75^\circ, 150^\circ)$

two test images is the mean intensity value (Figure 8.12 (a) and (d)); for shallower lighting the image looks darker. The shadowed areas in the image with zenith  $75^\circ$  are the same areas that are shadowed in the image with zenith  $45^\circ$ . Shadows are estimated for zenith  $45^\circ$  and encoded in the albedo map. Then they are added to the synthetic images rendered with any zenith angle. Since the shadowed areas remain almost the same for different light zeniths, the prediction of shadows is reasonably accurate. Consequently, shadows are better predicted for our directional textures than for the isotropic surfaces (see Figure 8.9).

The surface grain orientation is  $90^\circ$  for the *sand-rippled surface*. We investigate the effect that varying the light azimuth has on the performance of the Linear method (Figure 8.13). The SRR drops when the light azimuth is parallel to the surface grain, i.e.  $\tau = 90^\circ$  and  $270^\circ$ . Even the estimation image is poorly predicted (SRR = 5 dB). Although the average accuracy is over 10 dB, the surface appearance is not accurately predicted in the grain direction.

The 3-light PS estimate does not rely on a linear surface reflectance. Does




Figure 8.13: *Sand-rippled surface*: Accuracy of prediction when rendering the Linear Photometric Stereo estimate

this technique give a better surface estimate? In Figure 8.14 we plot the accuracy of prediction against light azimuth for this method. The prediction of the three estimation images is extremely accurate ( $\text{SRR} > 25 \text{ dB}$ ); even when the azimuth direction is aligned to the surface grain. However, apart from the estimation image, the prediction is poor when the light azimuth and the surface grain are parallel. This technique estimates the surface after removing the non-linearities of the image. When non-linearities are dominant, the ratio between the surface image and the non-linear terms is almost one and the surface estimate is not accurate.

Neither algorithm can reliably be used for predicting the appearance of a unidirectional surface lit from the surface grain direction.

### 8.3 Summary and discussion

In this chapter we assessed the proposed surface recovery method on real data. Since ground truth data was not available for the real textures, the assessment was in terms of image prediction. The Optimal Linear Filter was not assessed because it relies on ground truth data for the filter's training. Due to the limited number of




Figure 8.14: *Sand-rippled surface*: Accuracy of prediction when rendering the 3-light Photometric Stereo estimate

textures considered for the assessment, the conclusions drawn in this chapter cannot be generalised. However, the results give us an idea of the scope of the proposed algorithms.

The 3-light Photometric Stereo technique performs better than the Linear Photometric Stereo for the test data. An increase on surface roughness or a violation of the Gaussian assumption degrade both estimations. The 3-light Photometric Stereo estimate is more robust to surface roughness. In contrast, the Linear Photometric Stereo is more robust to surface discontinuities. Furthermore, the Linear method estimates surfaces with a hybrid reflectance more accurately than the 3-light PS technique.

If our objective is image prediction the following points are relevant. The prediction is more accurate for lighting conditions similar to the lighting conditions of the estimation. The 3-light Photometric Stereo technique is more robust to changes in the light azimuth than the Linear technique. This is because the three estimation images evenly sample half of the azimuth range. In the case of the Linear estimation, the two necessary images are evenly placed covering just a quarter of the azimuth range.

In contrast, the Linear Photometric Stereo technique is more robust to increments in the light zenith. For raking light the surface image is more linear, thus it is closer to the rendered Linear estimate. However, shadowing effects increase for larger zenith angles, causing both estimations to deteriorate.

For the purposes of image prediction, both techniques are reasonably accurate at predicting the surface appearance. In contrast, the estimates are not always sufficiently accurate for the purpose of surface recovery. For most of the textures, the recovery accuracy of the 3-light Photometric Stereo is close to 10dB. However, the Linear Photometric Stereo method does not accurately estimate most of the considered surfaces. We need to improve the estimation, but not at the expense of requiring more information about the texture. We want a simple technique that uses the available information in a more efficient manner. This simple technique is the subject of Chapter 9.

## Chapter 9

# Improving the Surface Reconstruction

---

In Chapter 8 we assessed the recovery performance of the novel Linear Photometric Stereo technique and compared it to that of the Benchmark technique. We found that whilst the Benchmark estimation was sufficiently accurate<sup>1</sup> for recovering rough surfaces, the Linear estimate needed to be improved. We argued that an algorithm which uses the available intensity data in a more efficient manner could be used for this purpose. In this chapter we review regularisation techniques to find a suitable algorithm. A simple iterative scheme that does not over-smooth the estimation is presented. The surface recovery assessment is repeated after implementing the iterative process as part of the estimation. An improvement in re-lighting performance is measured when the iterative scheme is used together with the Linear estimator.

---

<sup>1</sup>Within the definition of this PhD. See Section 8.1.1.

## 9.1 Chapter organisation

This chapter starts by reviewing regularisation algorithms that have been applied to Shape from Shading problems. We propose an iterative scheme that improves the surface estimation. The differences between our simple iterative scheme and the reviewed regularisation methods are identified. Our recovery techniques are assessed when using the iterative scheme. We investigate the robustness of the improved estimates to surface discontinuities, roughness and for several topographic and reflectance functions. The assessment results are presented and discussed.

## 9.2 How could the Photometric estimate be improved?

In Section 4.1.2 we reviewed the so-called minimisation Shape from Shading techniques. These approaches are optimisation techniques implemented in an iterative manner. They rely on certain assumptions about the surface. The surface is estimated by minimising an energy term which is a function of the validity of the surface assumptions.

In this section we review a class of minimisation methods: regularisation techniques. This survey is intended to help us design an algorithm to improve the surface estimate. Since regularisation techniques aim to optimise the estimate, they seem a good choice for our purposes.

### 9.2.1 Regularisation schemes: literature review

A possible solution for single-image Shape from Shading is achieved using regularisation. To regularise an ill-posed problem means to convert it into a well-posed problem. Traditional regularisation relies on assumptions about the smoothness of the surface. This assumption ensures stability of the approximate solution, but can be insufficient to preserve the main qualitative characteristics of the function sought



[131].

In addition to conditions of smoothness, shape constraints on the solution may exist. Common shape constraints are non-negativity conditions, monotonicity, convexity, etc. These constraints have stabilising properties. Furthermore, the estimated approximate solution has the desired qualitative behaviour. However, they rely on strong assumptions about the surface shape.

Horn [18] proposed an iterative method which seeks a solution in the form of the surface gradient. The solution ensured that the image irradiance equation was satisfied (expression 4.1). He introduced a smoothness regularisation term to guarantee stability. This regularisation term penalised rapid changes in the gradient field. The error function to minimise had the form:

$$e = e_s + \alpha e_i \tag{9.1}$$

where  $e_s$  is the error due to departure from smoothness,  $e_i$  is the brightness error and  $\alpha$  is the regularisation parameter. This parameter should be large if brightness measurements are very accurate, and small if they are very noisy. Horn suggested a discrete iterative implementation of his minimisation scheme.

He later improved his iterative algorithm so that the exact solution was found even though a regularisation term was included [67]. The smoothness term was only used to stabilise the iterative scheme when it was far from the correct solution. It was then turned off as the solution was approached.

Hayakawa *et al.* [132] separated the surface estimate modification required because of the brightness error, from that required due to the smoothness constraint. This allowed the use a simple and fast algorithm for each modification. Using simulation, their method gave more accurate results with a faster convergence than Horn's algorithm. Although it was not proved to give the optimal solution.

Subbarao and Liu [133] used a regularisation scheme to recover shape from defocused data. In their algorithm, the similarity function,  $e_i$ , is the brightness error used by Horn. In contrast, the regularisation term,  $e_s$ , is a smoothness constraint

based on a local blur parameter. The performance of their approach was assessed on synthetic data. They compared their algorithm to a gradient descent type approach and a local search approach. They concluded that regularisation performed better.

Ferrie and Lagarde [134] proposed a reconstruction algorithm based on minimising the variation of surface curvature. They used the curvature constraint to stabilise the result of the local shading analysis. Results were presented on synthetic and real data, showing that the algorithm was robust to image noise and gave accurate estimates. This method was compared to Horn's and was proved to preserve peaks that were smoothed with Horn's estimation.

Worthington and Hancock [135] argued that schemes with a quadratic regulariser over-smooth the surface estimate. They proposed an algorithm where the smoothness constraint does not dominate the iterative process. Data-closeness was imposed as a hard constraint. They argued that to satisfy the image irradiance equation, the surface normals must lie on a cone of ambiguity. At each iteration the updated surface normal was free, as a result of the smoothing process, to lie outside the cone, but it was subsequently projected back onto the closest vector lying on the cone. This framework allowed them to use different regularisation terms and compare their performance on real and synthetic data. They concluded that curvature and gradient consistency provided better constraints than quadratic smoothness terms. They successfully applied this method to face recognition problems [136].

### 9.2.2 Discussion

In this chapter, we search for a similar approach: to optimise the Photometric estimate we should adopt some of the constraints reported in the literature. Surface constraints related to the surface smoothness, its topography, its integrability as well as brightness constraints have been proposed. We should avoid assumptions about the smoothness of our textures. We cannot use topographic constraints since no information about the local surface shape is available. Our recovery methods enforce integrability using a global algorithm; thus no integrability issues have to

be considered. Brightness constraints could be an option since surface intensity information is available. The same test images that are used to estimate the surface can also be used to improve this estimate. Constraints related to the image intensity are a good choice in our case.

## 9.3 A simple iterative method

There are three main differences between our iterative scheme and the ones discussed in Section 9.2.1. First, our method does not require a regularisation term to stabilise the solution. Since the system is not ill-posed and the scheme is stable [131]. Second, the brightness error accounts for all available images. Third, an accurate initial estimate is provided by the novel recovery methods.

A "two-step" recovery was first introduced by Ferrie and Lagarde [134]. They used a single-image Shape from Shading technique to obtain a first estimate of the surface. They then refined the initial estimate through a minimisation procedure.

### 9.3.1 Description of the algorithm

Horn's regularisation scheme [18] influenced the design of our iterative method to improve the Photometric estimate using the intensity information. Horn's scheme in the discrete case takes the form:

$$p_{xy}^{n+1} = \bar{p}_{xy}^n + \alpha(I_{xy} - \hat{I}(p_{xy}^n, q_{xy}^n)) \overline{\frac{\partial \hat{I}}{\partial p}} \quad (9.2)$$

$$q_{xy}^{n+1} = \bar{q}_{xy}^n + \alpha(I_{xy} - \hat{I}(p_{xy}^n, q_{xy}^n)) \overline{\frac{\partial \hat{I}}{\partial q}} \quad (9.3)$$

where  $I$  is the real image and  $\hat{I}$  is the image computed using expression 3.4. The new values of  $p$  and  $q$  at each grid point are obtained using the old values of  $p$  and  $q$  in evaluating  $\hat{I}(p, q)$ ,  $\partial \hat{I} / \partial p$  and  $\partial \hat{I} / \partial q$ .  $\bar{p}^n$  and  $\bar{q}^n$  are the local averages of  $p$  and  $q$ ; they are the regularisation terms and account for departure from smoothness.

For our method several images are considered to compute the brightness error. Horn [18] suggested an expression for the brightness error that accounted for more than one image:

$$e_i = \sum_{j=1}^n \alpha_j \sum_x \sum_y (I_j(x, y) - \hat{I}_j(x, y))^2 \quad (9.4)$$

where  $I_j$  is the brightness measured in the  $j^{th}$  image and  $\hat{I}_j$  is the corresponding reflectance map. The parameter  $\alpha_j$  weights the errors due to the different images. They are equal if the images are equally reliable.

We propose an algorithm that omits Horn's regularisation term and uses equation 9.4 for the brightness error. The iteration is similarly carried out in terms of improving the gradient estimates. The resulting equations are 9.5 and 9.6.

$$p_{xy}^{n+1} = p_{xy}^n + \sum_{j=1}^n \alpha_j (I_{j;xy} - \hat{I}_j(p_{xy}^n, q_{xy}^n)) \frac{\partial \hat{I}_j}{\partial p} \quad (9.5)$$

$$q_{xy}^{n+1} = q_{xy}^n + \sum_{j=1}^n \alpha_j (I_{j;xy} - \hat{I}_j(p_{xy}^n, q_{xy}^n)) \frac{\partial \hat{I}_j}{\partial q} \quad (9.6)$$

The new values of  $p$  and  $q$  are not computed using the *local averages* of the old values of  $p$  and  $q$  as in Horn's algorithm. No regularisation is necessary, thus no smoothness constraints are imposed. We assume Lambertian reflection to compute  $R(p, q)$ ,  $\partial R/\partial p$  and  $\partial R/\partial q$ , thus:

$$\frac{\partial \hat{I}}{\partial p} = -\frac{\cos \tau \sin \sigma}{(p^2 + q^2 + 1)^{1/2}} - \frac{(-p \cos \tau \sin \sigma - q \sin \tau \sin \sigma + \cos \sigma)}{(p^2 + q^2 + 1)^{3/2}} \quad (9.7)$$

$$\frac{\partial \hat{I}}{\partial q} = -\frac{\sin \tau \sin \sigma}{(p^2 + q^2 + 1)^{1/2}} - \frac{(-p \cos \tau \sin \sigma - q \sin \tau \sin \sigma + \cos \sigma)}{(p^2 + q^2 + 1)^{3/2}} \quad (9.8)$$

$\hat{I}$  is computed using equation 3.4 for both recovery techniques. Thus no linearity of the surface reflectance is assumed when optimising the LPS estimate.

Once the system converges to the optimal solution, the gradient fields are integrated into the surface estimate. We use a global algorithm that enforces integrability (see Section 4.5.2.).

### 9.3.2 When to stop iterating?

This is a natural question that arises in most iterative processes. We want the gradient of the solution to approach that of the surface, but we do not know what the underlying surface is. Common quantities to test closeness to the correct solution are [67]:

- The brightness error should be small. This error is not a useful stopping criteria, since it becomes small just after a few iterations.
- Rate of change of the solution with iterations

$$\int \int \left( \frac{dp}{dt} \right)^2 + \left( \frac{dq}{dt} \right)^2 dx dy \quad (9.9)$$

We can keep track of this rate of change and stop iterating when it becomes small. In most cases it helps to continue for a while after expression 9.9 stops changing rapidly since the solution often continues to adjust a bit.

- Departure from smoothness errors are common stopping criteria. However they are not suitable in our case, since we do not assume a smooth surface.
- Measures of lack of integrability are not a good choice. Integrability is considered once the optimal gradient estimate is found and not during the iteration.

The quantity that we use to stop iterating is the global rate of change of the brightness error:

$$\frac{d[\sum_{j=1}^n \alpha_j \sum_x \sum_y (E_j(x, y) - R_j(p, q))^2]}{dt} \quad (9.10)$$

We should note that all the estimation images are taken into account when computing the brightness error. The rate of change of this error rapidly becomes small after few iterations (typically 7 or 8). Besides we do not stop iterating since the solution continues to adjust. We compromise computational cost for a more accurate estimation; the total number of iterations doubles the initial quantity.

### 9.3.3 Summary

In this section we give a brief description of the recovery algorithm:

1. Each of the proposed Photometric Stereo algorithms is applied to real images to obtain the surface estimate.
2. The Photometric estimate is used as an initial estimate for the iterative scheme.
3. The brightness error is computed (equation 9.4): the estimate is rendered with Lambert's law and compared to the real images.
4. The new surface gradient is estimated using equations 9.5 and 9.6.
5. The rate of change of the brightness error is evaluated (equation 9.10).
6. The process is repeated from step 3 until the expression 9.10 is steadily small.

The two phases of the surface recovery are: first the Photometric Stereo estimation (step 1) and second the iterative scheme (steps 2 to 6). We refer to the new recovery algorithms as the Iterative Linear Photometric Stereo (ILPS) and the Iterative 3-light Photometric Stereo (I3PS).

## 9.4 Assessment of the iterative recovery through rendering

### 9.4.1 Experimental method

#### Description

In this section, we apply two-step recovery to real data. We assess the recovery in terms of image prediction. The surface estimate is rendered for lighting conditions *different* from the ones used for the estimation. These rendered estimates are compared to the test images of the surface, and the SRR of equation 8.2 is computed.

The assessment of the Iterative Photometric Stereo recovery methods is four-fold:

1. We evaluate whether the estimation improvement due to the iterative process compensates for the extra computation.
2. We assess the recovery performance on real textures.
3. We investigate the robustness of our methods to four surface characteristics: surface roughness, discontinuities, topography and reflectance.
4. We test our algorithms on the so-called special surfaces, i.e. textures that violate the surface assumptions.

## Data

The two-step recovery algorithms are applied to the real textures described in Chapter 2. Seventeen real surfaces that roughly comply with the surface assumptions are considered in the assessment. Four textures that severely violate the assumptions are also estimated.

## Rendering shadows

The estimated surface is rendered so that its appearance is closest to the test image. Since the model for shadows is not realistic, we should avoid shallow illumination. Nevertheless, the non-Lambertian textures are rendered for  $75^\circ$ , i.e. the anaglypta surfaces and the sand paper (see Section 2.4.3). For the rest of the data set, large zenith angles are not considered.

## Constant albedo

For the surface recovery assessment carried out in Chapter 8, the estimated surface albedo was used for rendering. We argued that the albedo should only be used as part of the assessment procedure, and not as part of the recovery. This distinction is important since the albedo can only be estimated with the 3-light PS technique. Therefore, it should not be used to improve the Linear estimate.

The iterative process is part of the estimation. To compute the brightness error

(step 3 in Section 9.3.3), the rendered surface estimate is compared to the test image. This rendering process is part of the estimation procedure, thus the albedo is not used with the Linear estimate. However, for the assessment we measure the accuracy of prediction and the albedo is used for rendering both Photometric estimates.


## 9.4.2 Experimental results

### Is it worth the extra computation?

The recovery is assessed in terms of accuracy of prediction: we measure the brightness error of images that are not used for the estimation. On the other hand, the iterative scheme minimises the brightness error of the estimation images. It is possible that the surface estimate after convergence gives a smaller brightness error of the estimation images but larger for any other lighting condition. This is the case for the Iterative 3-light Photometric Stereo scheme. The surface estimate is refined to minimise the brightness error, but the accuracy of prediction is not improved or only slightly improved. The mean brightness error averaged across all textures is 19.53 dB for the 3PS and 20.41 dB for the I3PS (Appendix B: Section B.3). However, the accuracy of prediction is not so clearly improved (Figure 9.1). Only four out of seventeen textures gave a better accuracy of prediction, whilst two gave the same result. In most applications, it will be not worth using the iterative scheme to improve the 3-light Photometric Stereo estimate.

In contrast, the Iterative Linear estimate proves to be more accurate than the Linear estimate for most of the test textures (Figure 9.2). The average improvement in terms of image prediction is 1.39 dB. The accuracy of image prediction before and after iteration for the 3-light and the Linear estimates is included in Appendix B (Section B.3).





\*The 3-light estimate of the texture *6a* gives zero brightness error, so the iterative process cannot be applied

Figure 9.1: Accuracy of image prediction for the 3-light PS and the Iterative 3-light PS techniques




Figure 9.2: Accuracy of image prediction for the Linear PS and the Iterative Linear PS techniques




Figure 9.3: Accuracy of recovery for the Iterative Linear and the 3-light PS techniques

### Assessment of surface estimation through rendering

Figure 9.3 shows the accuracy of recovery of the real textures. The performance of Iterative Linear PS technique is very similar to that of the 3-light PS algorithm. The improvement of the Iterative Linear estimate has made its accuracy comparable to that of the 3-light PS estimate.

Ten out of the seventeen considered textures were accurately estimated with both methods, i.e. the recovery accuracy is larger than 10dB. The textures that have not been properly recovered do not comply with the surface assumptions. The *anaglypta surfaces* (Figure 2.6 (e), (f) and (g)) are not Lambertian. The *fracture* of Figure 2.5 (a) is very rough and it has heavily shadowed images. The very sparsely *deposited* surfaces of Figure 2.5 (i), (j) and (k) cannot be considered Gaussian.

Appendix C includes the 3D plots of the heightmaps estimated with both techniques for each of these textures. Both estimates are also rendered for two different lighting set-ups and compared to the real images.

Texture	Figure	ILPS	3PS
<i>Rough fracture</i>	2.5 (a)	8.69	8.52
<i>Moderate roughness fracture</i>	2.5 (b)	13.19	13.13
<i>Gentle roughness fracture</i>	2.5 (c)	12.56	12.71
<i>Smooth cured fracture</i>	2.5 (d)	11.89	13.37
<i>Smooth patterned fracture</i>	2.5 (e)	11.81	12.84

Table 9.1: Average accuracy of recovery for fractures [dB]




Figure 9.4: *Gentle roughness fracture* imaged for  $\sigma = 45^\circ$  and  $\tau = 300^\circ$

### Robustness to surface roughness

We compare the robustness of both estimators to surface roughness. Table 9.1 shows that both techniques are similarly affected by an increase in roughness. The accuracy of prediction for the Iterative Linear estimate is analogous to that of the 3-light PS. Neither method is accurate at predicting the appearance of the roughest surfaces (for instance the *rough fracture*).

We study the recovery of the *gentle roughness fracture* as an example of an accurate estimation. Figure 9.4 (a) is the test image of the fracture for a zenith of  $45^\circ$  and an azimuth of  $300^\circ$ . The rendered 3-light PS estimate and the rendered Iterative Linear estimate are the images of Figure 9.4 (b) and (c). The SRR is 12.80dB for the 3-light estimate and 13.24dB for the Iterative Linear estimate. Both synthetic images are very close to the test image, the main difference is due to shadowing.

Figure 9.5 (a) is a section of the test image: an area of 50x50 pixels from top left

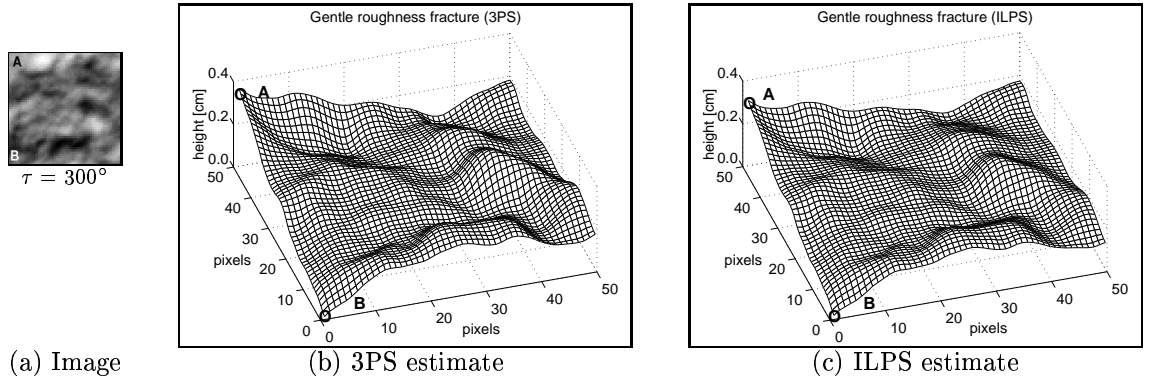


Figure 9.5: Estimated *gentle roughness fracture* surface

Texture	Figure	ILPS	3PS
<i>Moderate rough deposited fracture</i>	<i>2.5 (f)</i>	10.23	10.16
<i>Rough deposited fracture</i>	<i>2.5 (g)</i>	10.09	10.36
<i>Heavy deposited surface</i>	<i>2.5 (h)</i>	10.21	9.92
<i>Medium deposited surface</i>	<i>2.5 (i)</i>	9.05	9.14
<i>Light deposited surface</i>	<i>2.5 (j)</i>	7.81	8.67
<i>Sparse deposited surface</i>	<i>2.5 (k)</i>	7.15	7.04

Table 9.2: Average accuracy of recovery for deposited surfaces [dB]

corner. Figure 9.5 (b) and (c) are the heightmaps for that image section estimated with the 3-light and the Iterative Linear techniques. The estimated heightmaps are very similar to each other. Considering the high accuracy of image prediction for this fracture, the estimates appear to closely resemble the surface topography.

### Robustness to surface discontinuities

We argued that as the amount of deposit on a globally flat surface increases, the height distribution is closer to a Gaussian. A Gaussian surface is likely to be continuous. Table 9.2 shows that the proposed recovery techniques are similarly robust to violation of the Gaussian assumption. For sparsely deposited surfaces, the accuracy of prediction is not over the 10-dB threshold.

As an example of an unsuccessful recovery, we plot the estimated heightmaps




Figure 9.6: *Sparsely deposited texture* imaged for  $\sigma = 45^\circ$  and  $\tau = 300^\circ$

and the prediction of one of the images for the *sparsely deposited texture*. Figure 9.6 (a) is the test image of the texture for a light zenith of  $45^\circ$  and an azimuth of  $300^\circ$ . Figure 9.6 (b) and (c) are the 3-light PS estimate and the Iterative Linear estimate rendered for the same lighting conditions. The accuracy of predicting this test image is 4.04dB for the 3-light PS and 6.37dB for the Iterative Linear method. Although the SRR is well below the 10dB threshold, the synthetic images are fairly close to the test image. The synthetic images appear blurred compared to the test image; this is main source of error.

We recover an area of 50x50 pixels which corresponds to the left top corner of the image. Figure 9.7 (a) is the considered image section. Figure 9.7 (b) and (c) are the heightmaps estimated with the 3-light PS technique and the Iterative Linear method. Although the accuracy of prediction for this texture is the lowest of the whole set (surface  $5k$  in the graph of Figure 9.3), the two estimates are still consistent with each other. Furthermore, they seem to properly follow the topographic undulations that we perceive on the surface image.

### Surface estimation for several topographic and reflectance functions

Table 9.3 shows the accuracy of predicting the appearance of textures that differ in their topography and reflectance. Each texture follows a different surface function, although the sand ripples show a similar pattern. Furthermore, the sand ripples are Lambertian whilst the anaglypta surfaces have hybrid reflectance.

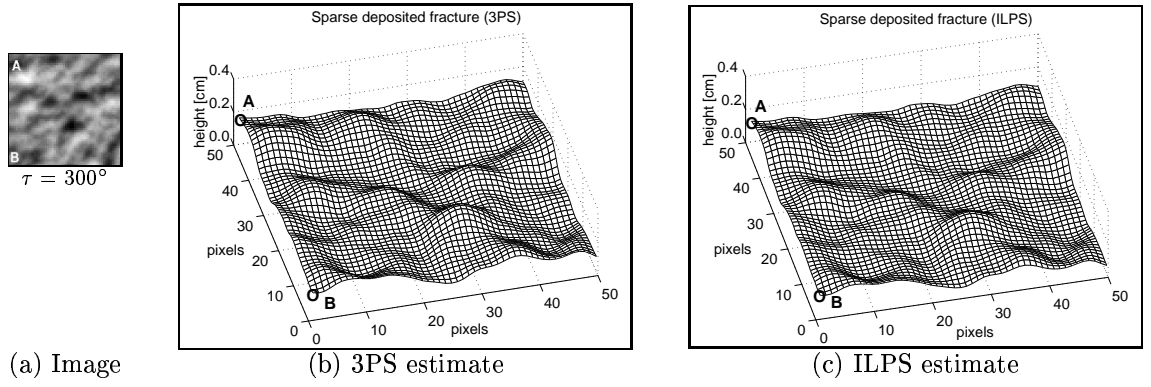


Figure 9.7: Estimated *sparsely deposited texture*

Texture	Figure	ILPS	3PS
<i>Sand-rippled surface</i>	<i>2.6 (a)</i>	9.58	11.75
<i>Sand-rippled surface (high freq)</i>	<i>2.6 (b)</i>	11.19	11.33
<i>Smoothed sand-rippled surface</i>	<i>2.6 (c)</i>	9.36	10.59
<i>Net anaglypta</i>	<i>2.6 (e)</i>	9.25	7.68
<i>Rippled anaglypta</i>	<i>2.6 (f)</i>	9.64	8.82
<i>Rice grain shaped anaglypta</i>	<i>2.6 (g)</i>	9.68	9.09

Table 9.3: Average accuracy of recovery for sand-rippled and anaglypta surfaces [dB]

As expected, both estimates are more accurate for the Lambertian sand ripples than for the anaglypta surfaces. However, the performance of the Iterative Linear technique is superior to that of the 3-light PS method for the anaglyptas. The Linear estimate is more robust to specularities than the 3-light PS estimate. We draw the same conclusion after the assessment through simulation in Chapter 7 and on real data in Chapter 8.

In contrast, the performance of the 3-light PS technique is better than that of the Iterative Linear method for the sand-ripples. It is worth noting that the rippled grain is aligned to the  $90^\circ$  direction. For light azimuths close to  $90^\circ$ , non-linearities become dominant and the Iterative Linear estimate deteriorates. Similarly, the Iterative Linear estimate of the *rippled anaglypta* would be more accurate for estimation images with a light azimuth different from  $90^\circ$ .

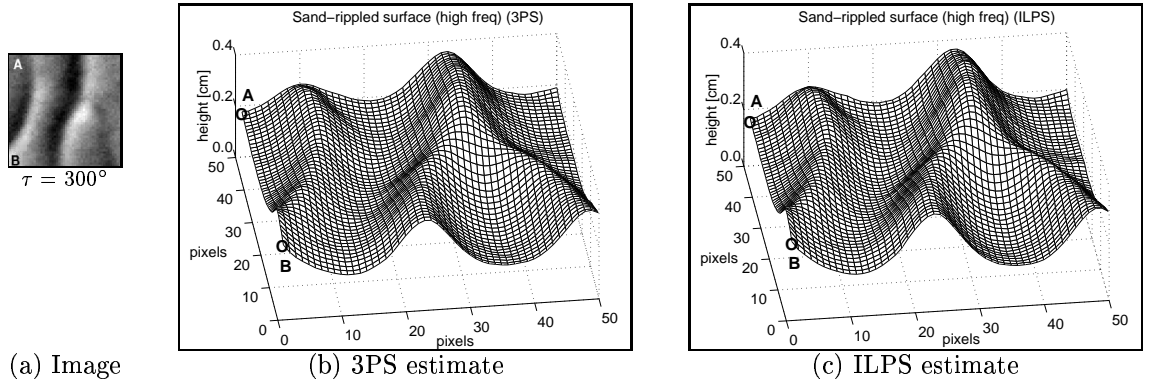


Figure 9.8: Estimated *high frequency sand ripple* surface

As an example we recover a 50x50 pixel section of the *high frequency sand ripple*. Figure 9.8 (a) is the image section which belongs to the top left corner of the image ( $\sigma = 45^\circ$  and  $\tau = 300^\circ$ ). Figure 9.8 (b) and (c) are the heightmaps estimated with the 3-light technique and the Iterative Linear method. Both estimates look identical and they seem to accurately follow the undulations of the surface image. We can easily distinguish the ripples in the estimated heightmaps.

### Surface estimation for special textures

In this section we apply our iterative recovery techniques to textures that severely violate the surface assumptions. In Chapter 8 we concluded that the recovery techniques were not suitable for recovering such surfaces. We here assess if the iterative process improves the estimation of the special textures. Table 9.4 shows the accuracy of image prediction for the recovery techniques before and after iteration. The surface estimate is not been improved with the iterative process for the test textures. Furthermore, for an inaccurate initial estimate, the iterative scheme is slow to converge to a solution. More than 30 iterations are required for some of the considered textures. For a larger regularisation parameter, the scheme becomes unstable. The iterative scheme is stable and fast in convergence when the initial estimate is not far from the solution.

Texture	Figure	3PS	I3PS	LPS	ILPS
<i>Irregular stripes anaglypta</i>	<i>2.6 (d)</i>	1.92	1.16	5.13	4.86
<i>Repetitive primitive surface</i>	<i>2.6 (i)</i>	5.76	-2.07	3.84	2.87
<i>Textile</i>	<i>2.5 (l)</i>	4.43	1.35	2.89	2.69
<i>Sand paper</i>	<i>2.6 (h)</i>	-1.72	-7.17	0.57	-0.41

Table 9.4: Average accuracy of recovery for special textures [dB]

## 9.5 Conclusions

In Chapter 8 we concluded that the Linear PS estimate was not accurate for most of our textures. In this chapter, we proposed a simple iterative scheme to improve the estimation of our recovery techniques. This iterative scheme minimised the brightness error of the images used for surface recovery. The estimates obtained with the proposed Photometric algorithms were used as initial estimates for the minimisation process. Although the brightness error for the iterative 3-light PS method was decreased, the image prediction was not improved. The extra computation was not efficient at improving the surface estimate. In contrast, the Iterative Linear PS method performed better than the Linear PS method. This improvement in the Iterative Linear technique made its performance comparable to that of the 3-light PS. Importantly we found that the iterative technique is stable and fast in convergence partly due to an accurate initial estimate.

We conclude that both recovery techniques give accurate estimates of rough surfaces. Their performance is similar in terms of accuracy of image prediction. The Iterative Linear PS requires less data for the recovery than the 3-light PS technique. In contrast, it requires more computation: the optimisation process. Furthermore, it is restricted to constant albedo surfaces, whilst the 3-light PS recovers the surface albedo as well as the surface gradient.



# Chapter 10

## Conclusions

---

### 10.1 Summary

This thesis has investigated the recovery of the topography of textured surfaces from their appearance. The texture's appearance depends on its topography, on its reflectance characteristics and on the lighting conditions. In Chapter 2 we discussed models of surface topography. In Chapter 3 we considered models of reflectance. In Chapter 6 we investigated the optimal lighting conditions. Chapter 4 reviewed the techniques of the literature that estimate shape from intensity. Chapter 5 presented the novel work on surface recovery. Chapters 7, 8 and 9 provided an assessment of the novel recovery techniques on synthetic and real data.

In Chapter 2, we defined the surfaces in terms of their topography. Our test surfaces are globally flat rough textures with low slope angles. They were assumed to be mathematically smooth continuous functions with conservative gradient fields. We presented the synthetic models that were used in the simulations throughout the thesis. Then we described our real textures and assessed how well they met our assumptions.

Our objective is to estimate the surface topography from its image. We therefore need to understand the transfer function from surface to image. In Chapter 3 we reviewed reflectance models. Lambert's law and Phong's model were thoroughly described since they were used to render diffuse and specular surfaces in our simulation. Kube's linear approximation to Lambert's law was reviewed because it can be inverted and is therefore a useful tool for surface estimation. We tested Kube's linear approximation on synthetic data. The conditions for the validity of the linear

model were extended to anisotropic surfaces. Kube’s model does not account for specular reflection. We analysed the contribution of Phong’s specular term to the total reflection using simulation. Phong’s model parameterises specular reflection using two variables: an exponent that controls the width of the specular lobe, and a scaling factor that controls its strength. In the case of textured surfaces lit from large zeniths, the scaling factor is the parameter that most significantly affects the surface appearance.

In Chapter 4 we reviewed existing surface recovery techniques. First, Shape from Shading techniques were surveyed. We argued that single-image Shape from Shading methods often rely on the assumption that the surface is smooth in order to estimate its gradient. In contrast, Photometric Stereo algorithms can be applied to rough surfaces. Since our objective is to recover a height field, the estimated gradient fields had to be integrated. Several integration methods were reviewed and we concluded that global integration was the most appropriate for our purposes. We introduced an algorithm that combined a successful 3-light Photometric Stereo technique with a well-known frequency integration method. We referred to this algorithm as the Benchmark Photometric Stereo technique.

Having reviewed previous work, in Chapter 5 we presented our contribution. We sought to recover rough surfaces more efficiently than with the Benchmark technique. We chose to decrease the number of intensity images without compromising on recovery accuracy. We surveyed two-light Photometric Stereo algorithms and investigated the conditions for a unique solution. We found that we could avoid the requirement of smoothness assumptions about the surface by assuming a linear reflectance function. Two novel two-light recovery algorithms were then introduced: the Linear Photometric Stereo technique and the Optimal Linear Filter. The Linear Photometric Stereo technique relied on Kube’s linear approximation to the surface reflectance function. The Optimal Linear Filter was trained on the test surface and its image. The estimated linear reflectance function is optimal for the *training surface* in the least squares sense.

Chapter 6 was devoted to optimising the conditions for the implementation of the new recovery methods. The optimal lighting geometry was investigated by reviewing the literature and conducting simulations. We investigated the optimal position of lights concluding that perpendicular lighting directions gave most accurate estimates. In terms of light zenith angle, the optimal range was found to be between  $30^\circ$  and  $60^\circ$ . For a more vertical lighting, the linear term was small relative to noise and nonlinearities. For shallower angles, shadowing effects were dominant. Although two images were sufficient to recover the surface, we studied the improvement when using more lights. We searched for a compromise between number of images and accuracy of estimation. However, we concluded that the improvement was not sufficient to justify the extra expense.

In Chapter 7 we assessed our novel techniques on synthetic data. The performance of our algorithms was compared to that of the Benchmark technique. The effect of violating the surface assumptions was studied. We tested the robustness of our estimators to the type and roughness of the test surface as well as its reflectance. We also assessed our estimators sensitivity to shadowing effects and noise. The Benchmark Photometric Stereo technique was found to be more robust to surface roughness than the novel algorithms. Both the Linear Photometric Stereo and the Optimal Linear Filter were limited to surfaces whose roughness is  $p_{rms} \leq 0.35$ . The Optimal Linear Filter was found to be more robust to surface specularities. The Linear Photometric Stereo estimate as well as the Benchmark estimate were very sensitive to specularities. However, the Linear Photometric Stereo estimate was found to be more accurate at recovering diffuse surfaces which have some specular contribution. We concluded that the Benchmark technique was the most robust alternative to shadowing effects. All three algorithms performed well under severe levels of noise.

In Chapter 8, we identified some of the limitations of using simulation to test our algorithms and extended the assessment to real data. Since fully calibrated ground truth data was not available, the assessment was in terms of image prediction.

The performance of the Linear Photometric Stereo method was compared to that of the Benchmark technique. The Benchmark technique outperformed the Linear Photometric Stereo method. Furthermore, we found that, for most of the test textures, the Linear Photometric Stereo estimate was not sufficiently accurate. The surface recovery had to be improved. However, both estimators were able to predict the surface appearance with reasonable good quality. For the purpose of image prediction, both methods presented valid alternatives.

In Chapter 9 we proposed an iterative technique that improves the Photometric estimate by minimising the brightness error. This simple iterative method was found to improve the Linear estimate, but not the Benchmark estimation. The improvement of the Linear Photometric Stereo technique made its performance comparable to that of the Benchmark technique. We concluded that both techniques are suitable to recover rough surfaces.

## 10.2 Conclusions

This thesis has investigated the issue of recovering the 3D structure of rough surfaces. In the literature very few estimators can recover the shape of rough surfaces. A recovery technique was introduced. The novelty was the combination of two standard components into an algorithm that could be applied to rough surfaces. These two components were a successful 3-light Photometric Stereo technique and a well-known global frequency integration method.

We argued that a 2-light Photometric Stereo technique could be used to find a unique solution of the surface without relying on surface smoothness assumptions. We proposed two novel recovery algorithms which estimated the shape of rough surface from only two images: the Linear Photometric Stereo technique and the Optimal Linear Filter. These algorithms relied on assumptions about the linearity of the surface reflectance.

The ability of the three proposed recovery methods to estimate rough surfaces

was assessed using simulation. We concluded that the estimates were accurate for fairly rough textures ( $p_{rms}, q_{rms} \leq 0.35$ ) which had a nearly Lambertian reflectance. The robustness of our estimates to shadowing effects and noise was also demonstrated.

We carried out an assessment of the Linear and the 3-light Photometric Stereo techniques on real data. We concluded that whilst the 3-light Photometric Stereo estimate was sufficiently accurate, the Linear estimate needed to be improved. We then proposed an iterative technique that improved the Linear estimate by minimising the brightness error. Both the improved Linear recovery and the 3-light Photometric Stereo recovery were found to perform well on real rough surfaces.

# Appendix A

## Slope distribution for the real surfaces

---

To estimate the gradient fields of our real textures we use the Photometric Stereo technique presented in Section 4.5.1. Four images per texture are used for the estimation. The lighting conditions of the images correspond to the smallest considered zenith angle and azimuth angles in  $90^\circ$  increments. Once the surface gradient is recovered, the mean and variance of the gradient fields are calculated. The slope distributions are plotted together with a Gaussian distribution which shares the same mean and variance. Modelling data is a problem in minimisation in many directions. However, there exist special, more efficient, methods that are specific to modelling, such as least squares methods or maximum likelihood algorithms. The Gaussian fit to our data is not optimal in the least squares sense nor it is the most likely curve. However, we use it because it is a simple but good first approximation.

By analysing the shape of the gradient distributions we can infer several relevant characteristics of the textures. A globally flat surface has zero mean in both gradient distributions. A large variance translates into a large *rms slope*, i.e. a rougher surface. A surface that has a large number of facets with high slopes is likely to be discontinuous.

We should bear in mind that the plotted distributions are estimates. Thus they might contain errors due to noise or to the estimation process. The validity of the conclusions drawn from studying these distributions depends on the quality of the estimation.

Figures A.1, A.2 and A.3 show the slope distributions for the *fractures*, the



*deposits* and the *textile*. These textures seem to follow a Gaussian distribution quite closely. Exception are the  $q$  derivative distribution of the *rough fracture* (Figure A.1 a) and the *smooth cured patterned fracture* (Figure A.2 a). The rest of the distributions are not perfectly Gaussian but they are clearly bell shaped. Their mean value is zero or approximate zero.

The ideal curve is more accurate estimating the number of facets with high slopes. There is no pattern for an over-estimation or under-estimation of the number of facets with low slopes.



The *roughness* of surface determines the distribution shape; rougher surfaces have wider distributions. Examples are the *rough fracture*, the *moderate roughness fracture*, the *moderate roughness deposit fracture*, the *rough deposit fracture*, the and *heavy deposited surface*.

It is worth noting that the *textile* is very well modelled by a Gaussian. This surface is considered among the special textures because it is rough and discontinuous. However the slope distribution does not show the characteristic shape of a rough surface. The reason is that the estimation process is not accurate for this surface.



We argued that as the amount of deposit decreases, the Gaussian assumption is less valid [33]. However, the considered *deposit surfaces* seem equally Gaussian. The ideal curve that we are comparing our estimates with is not the optimal Gaussian fit. For the optimal Gaussian fit and performing a chi square test, we would verify that the *sparsely deposited surfaces* are phase rich.





(a) Rough fracture (Figure 2.5 (a))



(b) Moderate roughness fracture (Figure 2.5 (b))





(c) Gentle roughness fracture (Figure 2.5 (c))





(d) Smooth cured fracture (Figure 2.5 (d))

Figure A.1: Slope distributions of real surfaces (real and fitted Gaussian)









(a) Smooth cured fracture with patterns (Figure 2.5 (e))



(b) Moderate roughness deposit fracture (Figure 2.5 (f))




(c) Rough deposit fracture (Figure 2.5 (g))





(d) Heavy deposited surface (Figure 2.5 (h))



Figure A.2: Slope distributions of real surfaces (real and fitted Gaussian)





(a) Medium deposited surface (Figure 2.5 (i))



(b) Light deposited surface (Figure 2.5 (j))



(c) Sparse deposited surface (Figure 2.5 (k))



(d) Textile (Figure 2.5 (l))




Figure A.3: Slope distributions of real surfaces (real and fitted Gaussian)

The second set of textures considered are the *sand ripples*, the *anaglyptas*, the *sand paper* and the *repetitive primitive surface*. Their slope distributions are plotted together with the ideal Gaussian in Figures A.4, A.5 and A.6.


These textures are phase rich, thus their gradient distributions are not Gaussian. Furthermore, the distributions of the *sand ripples* are not zero-mean.

For the *sand ripples* the surface grain is aligned to the  $90^\circ$  direction. The  $p$  derivative of the *low frequency sand ripple* captures the rising and falling of the ripples. Thus the  $p$  distribution shows more facets with a slope of  $26^\circ$  than flat facets. However, the less directional *sand ripples* do not model the ripples so clearly. As expected the  $q$  derivative distributions are more Gaussian.

The *anaglyptas* differ in their topographic functions. The *irregular stripes anaglypta*, the *net anaglypta* and the *ripples anaglypta* are rough surfaces showing wide distributions. In contrast, the *rice grain shape anaglypta* is comparatively smooth.


The *sand paper* is a rough surface which has some very specular facets. The estimated distribution do not show the surface roughness. The Photometric Stereo technique used to estimate the surface gradient is restricted to near Lambertian surfaces. The estimation error for this textures is large, and the estimated distribution is not approximate to the real distribution.

Finally, the *pearl barley* surface is a very rough discontinuous texture. All the surface discontinuities have been smooth out in the estimation process. Thus the slope distributions do not show a large number of facets with very high slope angles, close to  $-1$  or  $1$ . However the distributions correspond to a very rough surface.





(a) Low frequency sand ripple (Figure 2.6 (a))



Figure A.4: Slope distributions of real surfaces (real and fitted Gaussian)





(a) High frequency sand ripple (Figure 2.6 (b))



(b) Less directional sand ripple (Figure 2.6 (c))



(c) Irregular stripes anaglypta (Figure 2.6 (d))



(d) Net anaglypta (Figure 2.6 (e))










Figure A.5: Slope distributions of real surfaces (real and fitted Gaussian)



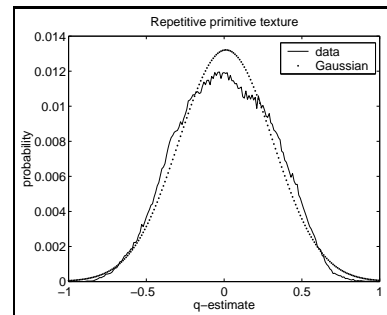

(a) Rippled anaglypta (Figure 2.6 (f))



(b) Rice grain shape anaglypta (Figure 2.6 (g))



(c) Sand paper (Figure 2.6 (h))



(d) Pearl barley (Figure 2.6 (i))

Figure A.6: Slope distributions of real surfaces (real and fitted Gaussian)

# Appendix B

## Result tables

---

### B.1 Should we use albedo for rendering the Linear estimate?

On the one hand, we should render the surface estimate using a model as realistic as possible. The albedo can be estimated with the 3-light PS techniques, thus it is available to us. On the other hand, estimation errors and noise contained in the albedo are added to the rendered Linear estimate. We here study how the accuracy of prediction is affected by considering the albedo for rendering.

The test textures — excluding the special surfaces — are estimated with the Linear PS method. The albedo map of these textures is estimated with the 3-light PS technique. We then render the Linear estimate with and without albedo and calculate the accuracy of prediction. In this experiment, we only consider the test images lit with the estimation zenith angle. There are twelve of these test images per texture. Two of them are used for recovering the surface, thus they are not taken into account to compute the accuracy of prediction.

Table B.1 shows the prediction accuracy measured when rendering with and without the surface albedo. The appearance of fourteen textures was better predicted when albedo was used for rendering the Linear estimate. The accuracy of prediction for the other three textures was degraded when considering the albedo. For these three textures the error contained in the albedo was large.

<b>Texture</b>	<b>without albedo</b>	<b>with albedo</b>
<i>Rough fracture</i>	8.46	8.61
<i>Moderate roughness fracture</i>	10.40	10.45
<i>Gentle roughness fracture</i>	8.95	8.95
<i>Smooth cured fracture</i>	8.08	8.39
<i>Smooth patterned fracture</i>	7.15	7.38
<i>Moder rough deposited fracture</i>	9.79	10.32
<i>Rough deposited fracture</i>	10.09	10.30
<i>Heavy deposited surface</i>	10.41	10.68
<i>Medium deposited surface</i>	9.97	10.11
<i>Light deposited surface</i>	8.79	10.17
<i>Sparse deposited surface</i>	8.24	8.08
<i>Sand-rippled surface</i>	9.19	10.31
<i>Sand-rippled surface (high freq)</i>	11.49	11.63
<i>Smoothed sand-rippled surface</i>	10.37	10.85
<i>Net anaglypta</i>	13.16	12.95
<i>Rippled anaglypta</i>	12.13	12.25
<i>Rice grain shaped anaglypta</i>	10.45	10.32

Table B.1: Accuracy of the predicting the real surface appearance using the Linear estimate [dB]

## B.2 Assessment of surface estimation through rendering

Table B.2 shows the accuracy of predicting the appearance of our real surfaces when rendering the 3-light PS and the Linear estimates. This table is complementary to the results discussed in Chapter 8, Section 8.1.2. The accuracy of prediction is been broken down for each zenith angle.

The 3-light PS performance for some textures deteriorates under raking light. Examples are the *rough fracture*, the two *deposited fractures* and all the *deposited surfaces*. These are surfaces whose roughness is due to rapid step changes of small amplitude. Furthermore, they are constant albedo surfaces and their estimated albedo maps mainly recover shadows and highlights. When the surface is lit with the estimation zenith, no shadows appear. However, for raking light the surface image is heavily shadowed. Therefore the estimated albedo contains none of the shadows that appear for raking lightings and the image is not properly predicted.

In contrast, the Linear estimate of these same textures is close to the test image when rendered for large zenith angles. The prediction is so accurate since the Linear estimator somehow accounts for shadows. The shadowing contained in the estimation images is "folded" into the Linear estimate. When the it is rendered using the albedo, shadows are again taken into consideration. Considering shadows twice makes the surface appearance closer to the test image for raking light.

For the rest of the data set, the accuracy of prediction is not so affected by changes in zenith. These other textures are fairly smooth surfaces except for some low frequency trends. Examples are the *sand ripples* and the *anaglyptas*. For these textures the areas that are shadowed remain the same for increasing zenith. However, a darker image is captured for raking light. The albedo map recovers the shadows of the estimation images. These shadows are then added to the estimates when rendering for any zenith. Since the shadowed areas are the same for any zenith, the prediction is steadily accurate for varying zenith.



	$\sigma = 45^\circ$		$\sigma = 60^\circ$		$\sigma = 75^\circ$	
<b>Texture</b>	<b>LPS</b>	<b>3PS</b>	<b>LPS</b>	<b>3PS</b>	<b>LPS</b>	<b>3PS</b>
<i>Rough fracture</i>	8.61	10.91	7.27	8.85	7.31	3.66
<i>Moderate roughness fracture</i>	10.45	15.24	10.68	14.83	10.10	7.73
<i>Gentle roughness fracture</i>	8.95	15.17	8.55	14.03	8.14	7.24
<i>Smooth cured fracture</i>	8.39	16.16	8.20	14.63	8.29	9.09
<i>Smooth patterned fracture</i>	7.38	17.19	6.68	12.78	6.46	8.41
<i>Moder rough deposited fracture</i>	10.32	12.74	8.98	10.77	8.82	5.41
<i>Rough deposited fracture</i>	10.30	13.20	9.13	10.81	9.15	3.57
<i>Heavy deposited surface</i>	10.68	12.87	8.81	10.19	9.01	3.24
<i>Medium deposited surface</i>	10.11	12.18	8.30	9.15	8.00	2.62
<i>Light deposited surface</i>	10.17	11.23	8.09	8.92	8.13	3.52
<i>Sparse deposited surface</i>	8.08	9.25	7.34	7.15	7.07	0.69
<i>Sand-rippled surface</i>	10.31	12.98	10.48	13.75	10.47	9.77
<i>Sand-rippled surface (high freq)</i>	11.63	12.29	12.32	13.43	12.15	10.92
<i>Smoothed sand-rippled surface</i>	10.85	11.81	11.18	12.31	11.09	9.21
<i>Net anaglypta</i>	—	—	12.95	9.31	9.28	8.37
<i>Rippled anaglypta</i>	—	—	12.26	9.68	10.46	10.39
<i>Rice grain shaped anaglypta</i>	—	—	10.32	8.62	10.76	11.72
<i>Irregular stripes anaglypta</i>	—	—	3.88	5.14	0.94	5.12
<i>Sand paper</i>	—	—	0.73	-0.66	0.54	-2.95
<i>Repetitive primitive surface</i>	4.29	7.42	4.11	5.95	2.12	2.12
<i>Textile</i>	2.96	6.04	3.32	4.33	—	—

Table B.2: Accuracy of the prediction of real surfaces appearance [dB]

## B.3 Reduction of the brightness error with the iteration process

Table B.3 shows the brightness error measured for both techniques before and after iteration. The brightness error is calculated as the average of the SRR for the estimation images of the texture (equation 8.1). For both techniques the appearance of the rendered estimates is closer to the estimation images after iteration. The Iterative Linear estimate gives a brightness error of less than 1% ( $\text{SRR} > 20\text{dB}$ ) for all textures. It shows an average improvement of almost 20dB with respect to the brightness error before iteration. In contrast, the SRR of the estimation images is not so drastically increased with the post-processing for the 3-light PS technique. It gives an average improvement of less than 1dB.

## B.4 Accuracy of prediction with the iteration process

Although the brightness error is been reduced, the accuracy of prediction might not be increased. This is the case of the 3-light PS technique (Table B.4). The estimator performs better for four of the seventeen textures.

In contrast, the performance the Iterative Linear technique is improved compared to that of the Linear technique. Twelve of the considered textures were more accurately predicted after iteration. Only the *anaglypta surfaces* and the *Light deposited surface* gave a less accurate prediction after iteration.


Figure B.1 plots the accuracy of recovery for the Iterative 3-light PS technique and the Iterative Linear method. The Iterative Linear estimator outperforms the Iterative 3-light PS estimator for nine of the seventeen surfaces. For the other eight surfaces, the accuracy of prediction is similar with both techniques.

Texture	LPS	ILPS	3PS	I3PS
<i>Rough fracture</i>	10.21	30.46	16.50	17.36
<i>Moderate roughness fracture</i>	10.62	29.68	18.75	20.58
<i>Gentle roughness fracture</i>	9.86	32.30	19.61	23.33
<i>Smooth cured fracture</i>	8.81	32.08	*0°, 180°	*0°, 90°, 180°
<i>Smooth patterned fracture</i>	6.99	30.64	*90°	*0°, 90°, 180°
<i>Moder rough deposited fracture</i>	11.70	34.31	18.42	20.19
<i>Rough deposited fracture</i>	11.96	29.50	22.30	12.48
<i>Heavy deposited surface</i>	11.97	33.57	18.79	19.26
<i>Medium deposited surface</i>	11.86	34.47	23.78	25.23
<i>Light deposited surface</i>	11.04	29.62	18.16	18.58
<i>Sparse deposited surface</i>	9.91	29.74	30.01	30.85
<i>Sand-rippled surface</i>	10.12	22.12	*0°, 90°, 180°	—
<i>Sand-rippled surface (high freq)</i>	10.90	22.17	27.78	30.11
<i>Smoothed sand-rippled surface</i>	11.56	*0°, 90°	29.50	32.25
<i>Net anaglypta</i>	14.11	37.39	9.76	12.71
<i>Rippled anaglypta</i>	13.15	31.89	10.21	11.16
<i>Rice grain shaped anaglypta</i>	11.79	30.11	9.88	11.65
* Brightness error is zero for the indicated images				
— Brightness error is zero for the three estimation images, the iteration cannot be performed				

Table B.3: Brightness error before and after iteration for the 3-light PS and the Linear PS [dB]

<i>Texture</i>	<i>LPS</i>	<i>ILPS</i>	<i>3PS</i>	<i>I3PS</i>
<i>Rough fracture</i>	7.22	8.69	8.52	8.53
<i>Moderate roughness fracture</i>	9.69	13.19	13.13	13.35
<i>Gentle roughness fracture</i>	8.00	12.56	12.71	11.72
<i>Smooth cured fracture</i>	7.60	11.89	13.37	12.20
<i>Smooth patterned fracture</i>	6.41	11.81	12.84	12.43
<i>Moderate rough deposited fracture</i>	8.79	11.70	10.16	9.99
<i>Rough deposited fracture</i>	8.86	10.09	10.36	10.37
<i>Heavy deposited surface</i>	8.86	10.71	9.92	10.14
<i>Medium deposited surface</i>	8.36	9.05	9.14	9.02
<i>Light deposited surface</i>	8.28	7.81	8.67	8.39
<i>Sparse deposited surface</i>	7.04	7.15	7.04	7.60
<i>Sand-rippled surface</i>	9.53	9.58	11.75	—
<i>Sand-rippled surface (high freq)</i>	11.01	11.19	11.33	10.61
<i>Smoothed sand-rippled surface</i>	10.11	9.36	10.59	9.89
<i>Net anaglypta</i>	10.01	9.25	7.68	8.03
<i>Rippled anaglypta</i>	10.34	9.64	8.82	7.95
<i>Rice grain shaped anaglypta</i>	9.68	9.68	9.09	8.27
— Since brightness error is zero for the three estimation images, the iterative process can not be carried out				

Table B.4: Accuracy of prediction before and after iteration for the Linear PS and the 3-light PS [dB]



\*The 3-light estimate of the texture *6a* gives zero brightness error, so the iterative process cannot be applied

Figure B.1: Accuracy of recovery for the Iterative 3-light PS and the Iterative Linear PS techniques

## B.5 Assessment of the iterative Linear estimate through rendering

Table B.5 shows the accuracy of the Iterative Linear technique. The measured accuracy is broken down for the considered zenith angles. In Section B.2 we argued that the Linear PS kept the same level of accuracy for increasing zenith. This is not the case after iteration; the Iterative Linear estimate shows a similar behaviour to the 3-light PS estimate for rough surfaces.

Texture	45°	60°	75°
<i>Rough fracture</i>	10.54	8.60	3.33
<i>Moderate roughness fracture</i>	14.41	14.37	6.90
<i>Gentle roughness fracture</i>	14.39	13.13	6.25
<i>Smooth cured fracture</i>	12.61	13.27	8.56
<i>Smooth patterned fracture</i>	14.51	11.52	8.94
<i>Moderate rough deposited fracture</i>	12.09	10.38	5.00
<i>Rough deposited fracture</i>	11.93	10.23	3.65
<i>Heavy deposited surface</i>	12.49	10.01	3.12
<i>Medium deposited surface</i>	11.35	8.64	2.91
<i>Light deposited surface</i>	9.31	7.87	3.01
<i>Sparse deposited surface</i>	8.64	7.11	-0.22
<i>Sand-rippled surface</i>	10.03	10.81	11.15
<i>Sand-rippled surface (phase-rich)</i>	11.75	12.58	10.71
<i>Smoothed sand-rippled surface</i>	10.41	10.04	9.31
<i>Net anaglypta</i>	—	11.34	9.06
<i>Rippled anaglypta</i>	—	10.36	10.64
<i>Rice grain shaped anaglypta</i>	—	10.18	10.87

Table B.5: Accuracy of prediction of the Iterative Linear PS technique against zenith

# Appendix C

## Reconstruction results

---

In this Appendix we show some reconstruction results for the seventeen real textures. The surface is estimated using both the 3-light PS technique and the Iterative Linear technique. These estimates are rendered for two different lighting conditions. We chose the lighting conditions so that either the zenith or the azimuth is close to the estimation angles. Thus we expect the rendered estimates to be fairly similar to the test images. For the *fractures*, the *deposits* and the *sand ripples*, the rendering lighting condition are:

1. First test image:
  - The zenith angle is the estimation zenith, i.e.  $45^\circ$ .
  - The azimuth angle is  $300^\circ$ , so it is very different from the estimation azimuth angles.
2. Second test image:
  - The zenith angle is  $60^\circ$ , thus it is different from the estimation zenith.
  - The azimuth angle is  $30^\circ$ , so it is close to the estimation azimuth angles.

For the *anaglypta surfaces* the rendering lighting condition are:

1. First test image:
  - The zenith angle is the estimation zenith, i.e.  $60^\circ$ .
  - The azimuth angle is  $300^\circ$ , therefore it is very different from the estimation azimuth angles.
2. Second test image:

- The zenith angle is  $75^\circ$ , thus it is different from the estimation zenith.
- The azimuth angle is  $30^\circ$ , so it is close to the estimation azimuth angles.

The rendered estimates are shown together with the corresponding test image. In general the main differences between the synthetic and the real images are due to shadows and highlights. Cast shadows are not modelled during the rendering of the estimates, thus they are not predicted. For very rough surfaces this effect is very distinct and it is a source of error. Highlights are not rendered either, thus the synthetic images differ from the test images whenever highlights occur.

A 3D plot of the estimates is also included per texture. Both the Iterative Linear estimate and the 3-light PS estimate are plotted. The 3D plots correspond to an area of 50x50 pixels at the top left corner of the surface sample. The image section that matches the surface area is magnified and shown together with the 3D plots. This section belongs to an image with zenith  $45^\circ$  and azimuth  $300^\circ$ . Including the magnified image section helps the reader to assess the recovery performance by looking at the 3D plots. It is important to note that the lighting for this image comes from the  $300^\circ$  direction.

In general both 3D estimates are very similar to each other. Furthermore, both seem to follow the undulations of the surface that we can infer from its image. The 3-light PS estimate shows more of the high frequency variation of the surface topography. When assessing and comparing our recovery techniques (Chapter 9) we observed that the 3-light PS method was more robust to roughness than the Iterative Linear method. This is the reason for an Iterative Linear estimate that looks like a slightly smoothed version of the 3-light estimate.



## Rough fracture

This texture is a good example of a very rough surface. In the test images cast shadows appear. These cast shadows are not rendered in the synthetic images.




Figure C.1: *Rough fracture*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

The two 3D estimates are very similar. However, the point **A** is displaced from one heightmap to the other. For the 3PS estimate **A** has a height of about 0.6cm. In the Iterative Linear estimate, the height value for **A** is about 0.5cm.

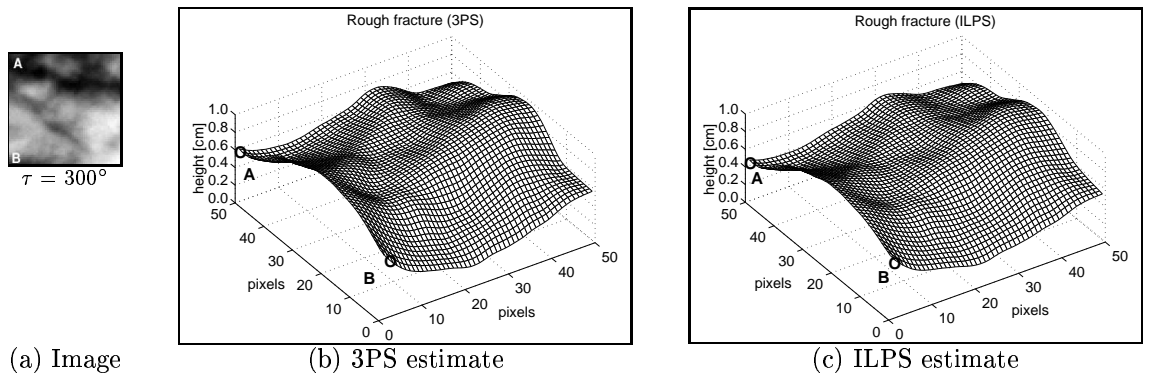


Figure C.2: Height recovery for the *rough fracture*

## Moderate roughness fracture

The synthetic images of the first row are a good quality prediction of the corresponding test image. However, the synthetic images of the second row fail to properly imitate the shadows. The effect of erroneously shadowed images is that rough surfaces appear smoother.




Figure C.3: *Moderate roughness fracture*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

This is a good example of how the Iterative Linear estimate looks like a flattered version of the 3-light estimate.

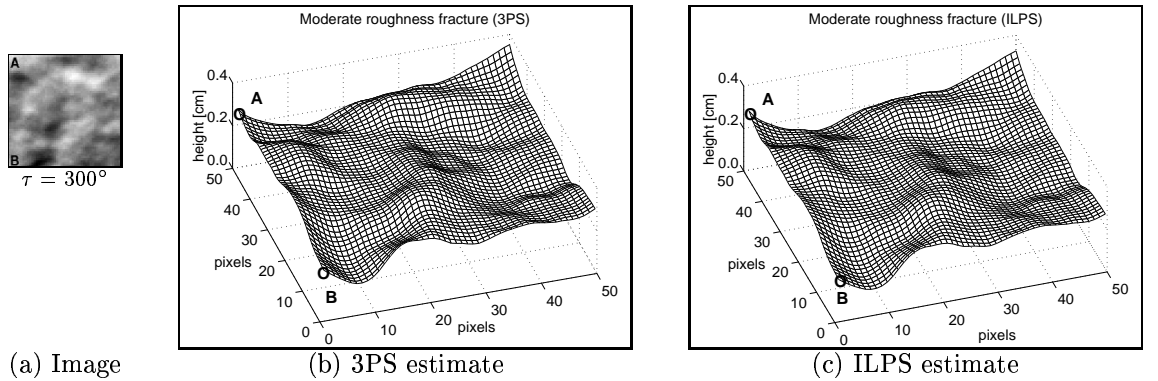


Figure C.4: Height recovery for the *Moderate roughness fracture*

## Gentle roughness fracture

The second test image is been quite accurately modelled. The rendered estimates show the roughness of the texture, although cast shadows are not predicted.




Figure C.5: *Gentle roughness fracture*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

Although both estimates are very similar, a difference can be appreciated in the height of point **A**. This values is 0.35cm for the 3-light PS estimate and 0.30cm for the Iterative Linear estimate.

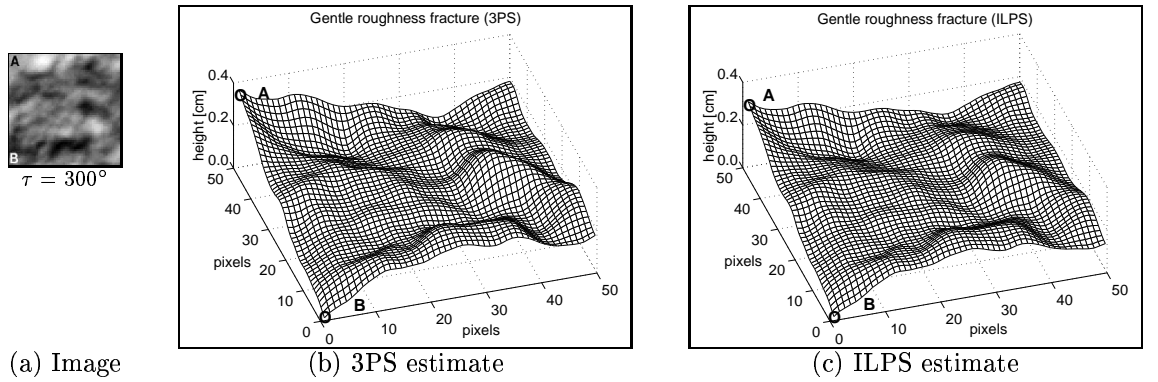


Figure C.6: Height recovery for the *gentle roughness fracture*

## Smooth cured fracture

Usually shadows are better predicted for the lighting condition of the second row images, i.e.  $(\sigma, \tau) = (60^\circ, 45^\circ)$ . The albedo, which mainly recovers shadows, is estimated for images with azimuths  $0^\circ$ ,  $90^\circ$  and  $180^\circ$ . The test images with azimuths close to the estimation angles show similar shadowing. This is the case of the second test image.




Figure C.7: *Smooth cured fracture*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

The Iterative Linear estimate looks smoother than the 3-light PS estimate.

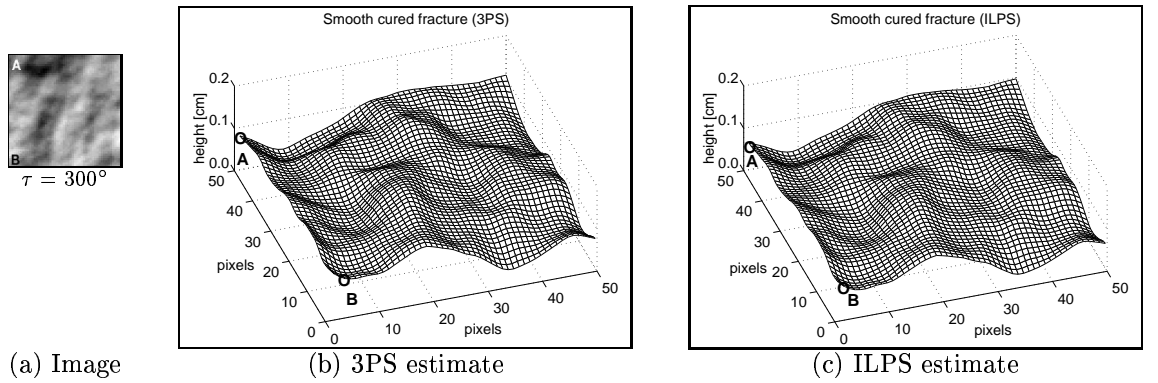


Figure C.8: Height recovery for the *smooth cured fracture*

## Smooth patterned fracture

The appearance of this smooth texture is accurately predicted. The effect of cast shadowing is not so distinct. It is worth noting the accurate prediction of the low frequency trend of the texture.




Figure C.9: *Smooth patterned fracture*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

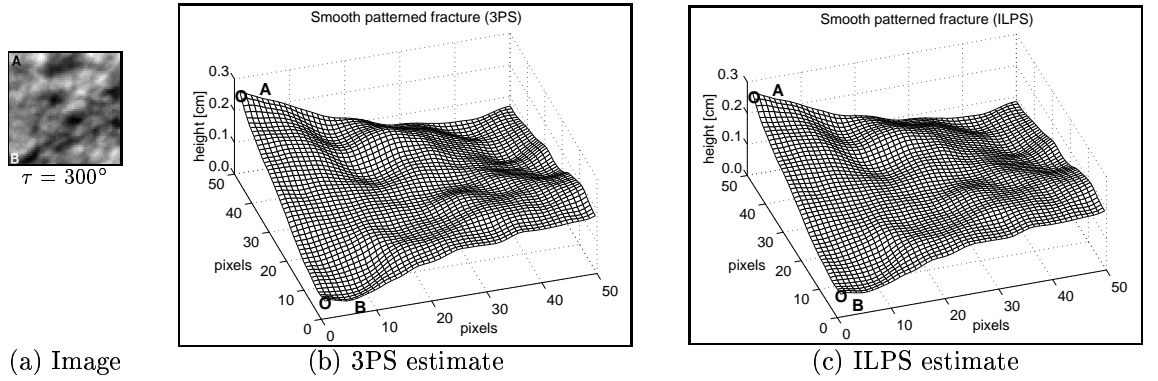


Figure C.10: Height recovery for the *smooth patterned fracture*

## Moderate roughness deposited fracture

This texture is fairly rough, however its appearance is properly predicted.




Figure C.11: *Moderate roughness deposited fracture*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

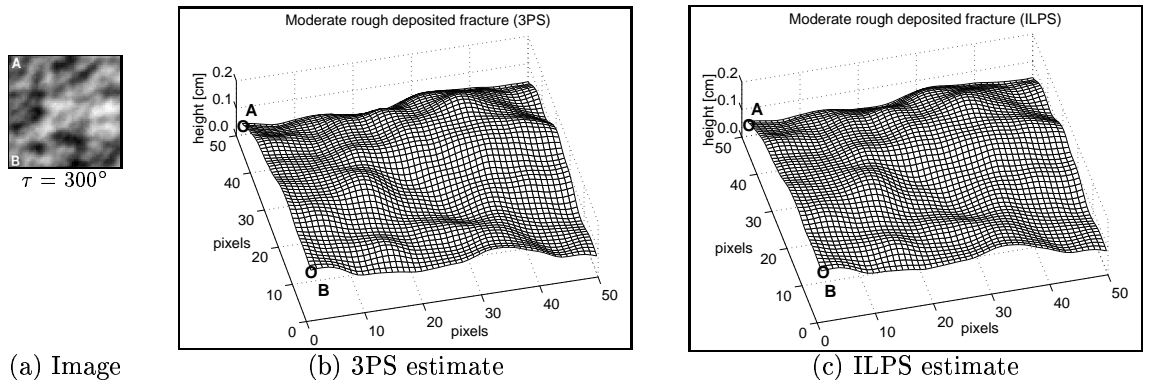


Figure C.12: Height recovery for the *moderate roughness deposited fracture*

## Rough deposited fracture

As the previous texture, this deposit fracture is a rough surface but still accurately recovered.




Figure C.13: *Rough deposited fracture*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

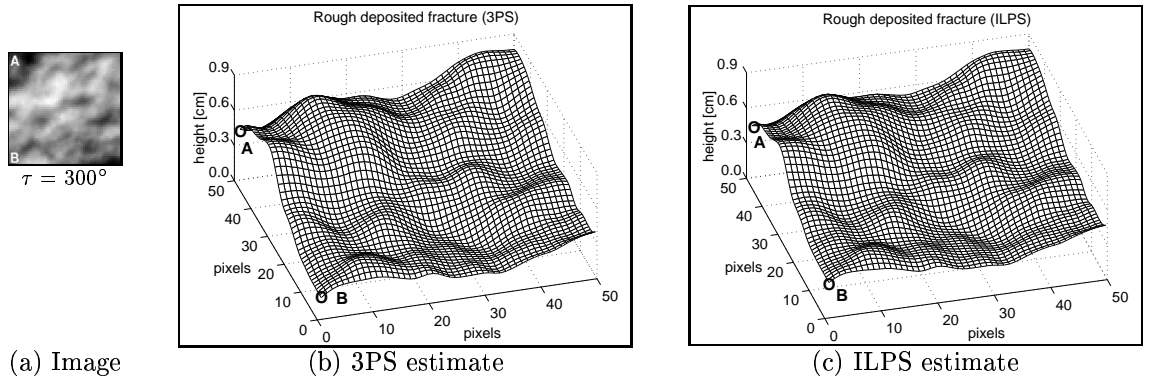


Figure C.14: Height recovery for the *Rough deposited fracture*



## Heavily deposited fracture

As the lighting is shallower, shadowing increases and the synthetic images differ from the test image.

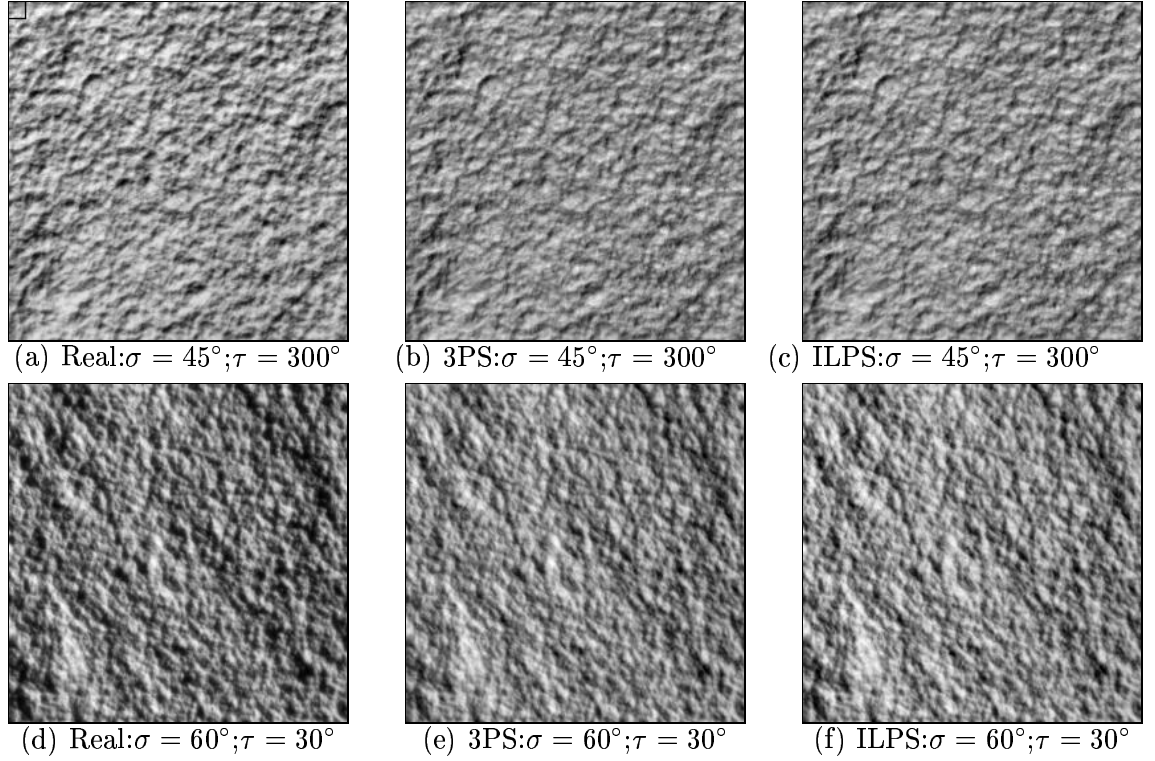


Figure C.15: *Heavily deposited fracture*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

Both estimates show the undulation that are seen in the image. However, high frequency variation of the surface topography due to the most recent deposits are not recovered.

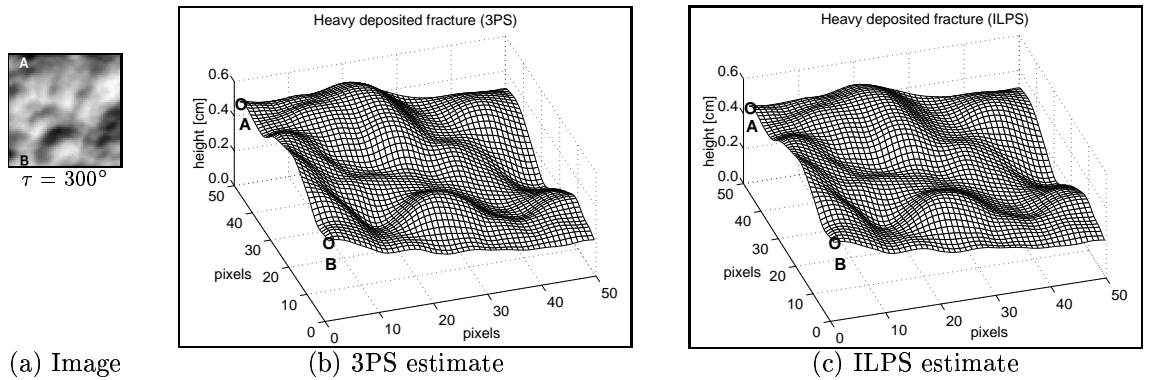


Figure C.16: Height recovery for the *heavily deposited fracture*



## Medium deposited fracture

The test images are very accurately predicted with both techniques.




Figure C.17: *Medium deposited fracture*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

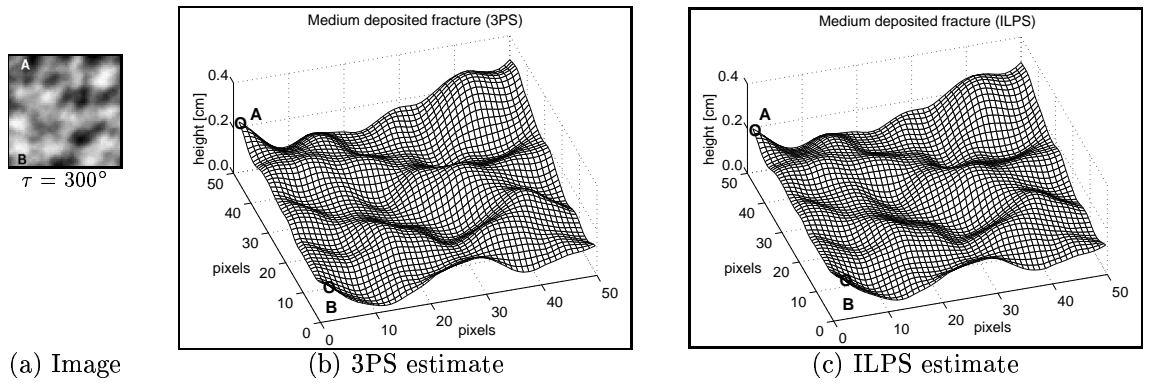


Figure C.18: Height recovery for the *medium deposited fracture*

## Light deposited fracture

The appearance of this texture is properly predicted with both estimates. The main difference is that the test image looks darker than the rendered estimates for shallow lighting.




Figure C.19: *Light deposited fracture*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

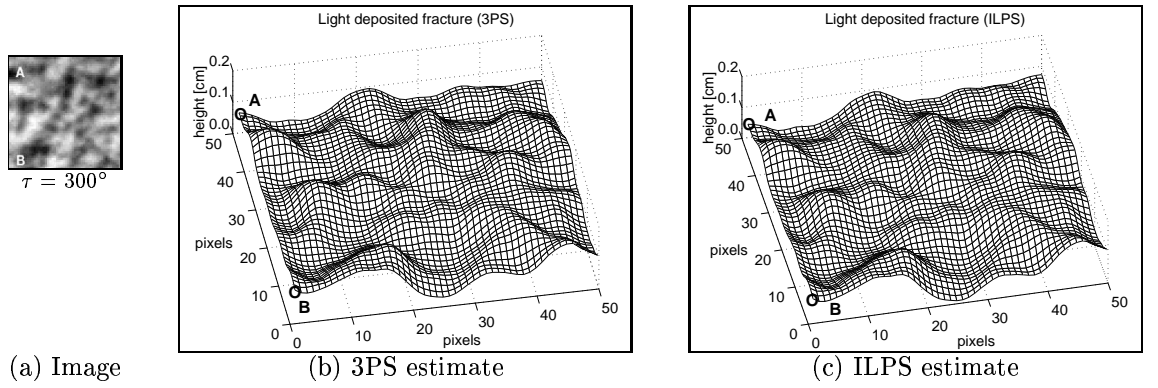


Figure C.20: Height recovery for the *light deposited fracture*

### Sparse deposited fracture




Figure C.21: *Sparse deposited fracture*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

The height undulations due to the most recent deposits are properly recovered.

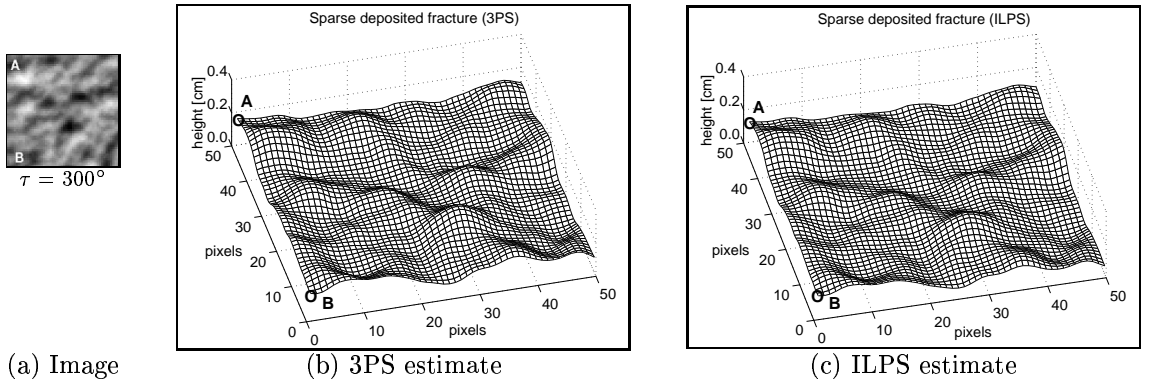


Figure C.22: Height recovery for the *sparse deposited fracture*

## Sand-rippled surface

The rendered Iterative Linear estimate is much lighter than the test image for shallow lighting. In contrast, the 3-light PS estimate shows the proper average intensity value.




Figure C.23: *Sand-rippled surface*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

The ripples are clearly appreciated in the 3D heightmaps.

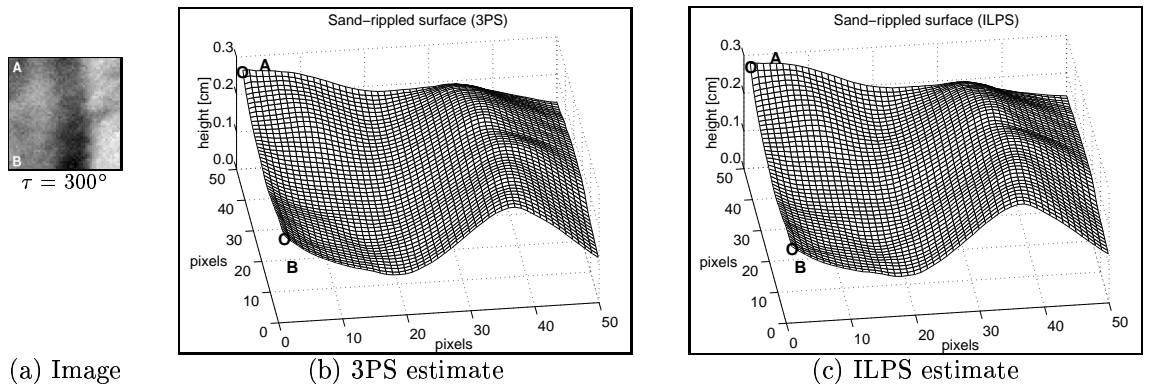


Figure C.24: Height recovery for the *sand-rippled surface*

## Sand-rippled surface (high frequency)

Both rendered estimates are very close to the test image even for shallow lighting.




Figure C.25: *Sand-rippled surface (high frequency)*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

The high frequency ripples are properly recovered.

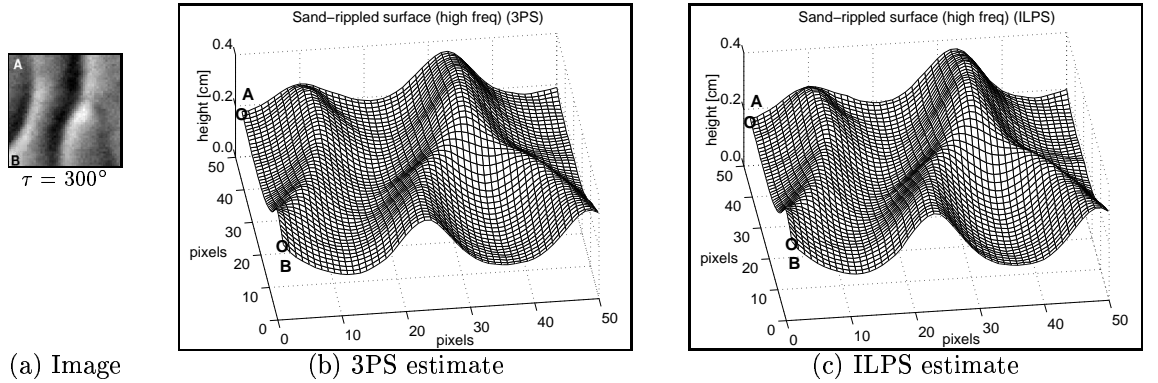


Figure C.26: Height recovery for the *sand-rippled surface (high frequency)*

### Smoothed sand-rippled surface




Figure C.27: *Smoothed sand-rippled surface*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (45^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (60^\circ, 30^\circ)$

The blurred effect that is observed in the images is also represented in the heightmap. Apart from the low frequency trend of the ripple, the heightmap shows some undulations of higher frequency and lower amplitude. These undulations are not observed in the other sand-ripples. They cause the sense of a blurred surface image.

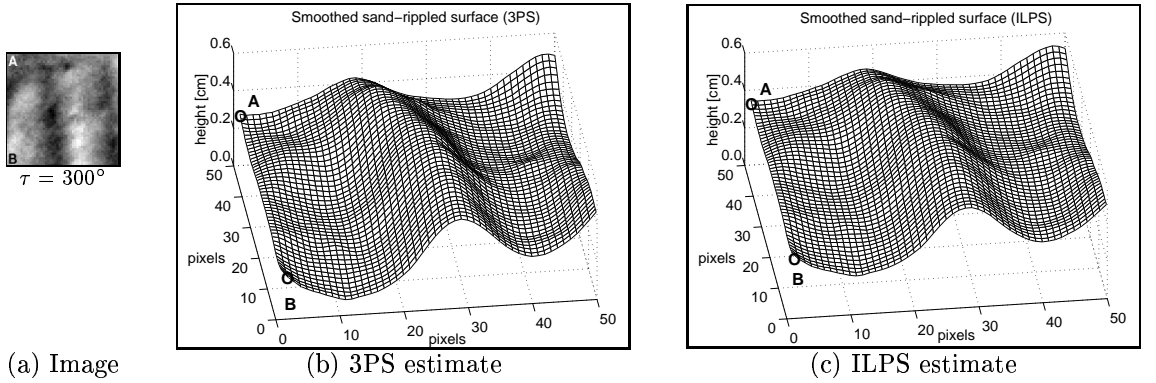


Figure C.28: Height recovery for the *smoothed sand-rippled surface*

## Net anaglypta

The synthetic images are very similar to the test images. The main difference is due to the mean intensity value. The effect is more noticeable for the rendered 3-light PS estimate.




Figure C.29: *Net anaglypta*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (60^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (75^\circ, 30^\circ)$

In general, the topographic information of both the sand ripples and the anaglyptas is easier to interpret. Thus the 3D estimates seem more accurate. We should note that this is not necessarily the case.




Figure C.30: Height recovery for the *net anaglypta*



## Rippled anaglypta

The anaglypta surfaces are not perfectly Lambertian. This can be perceived in the test images; some highlights are captured. Highlights are not modelled when rendering the surface estimates. However the synthetic images do not drastically differ from the test images.




Figure C.31: *Rippled anaglypta*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (60^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (75^\circ, 30^\circ)$

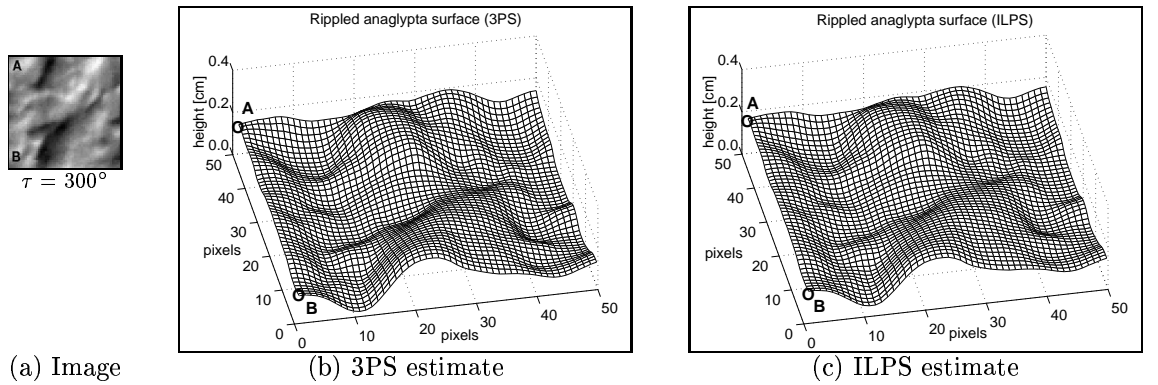


Figure C.32: Height recovery for the *rippled anaglypta*



## Rice grain shape anaglypta

This surface's appearance is properly predicted with the exception of the mean intensity value for the 3-light PS estimate rendered for raking light.




Figure C.33: *Rice grain shape anaglypta*: Real image (left); 3PS estimate (centre); ILPS estimate (right); top row:  $(\sigma, \tau) = (60^\circ, 300^\circ)$ ; bottom row:  $(\sigma, \tau) = (75^\circ, 30^\circ)$

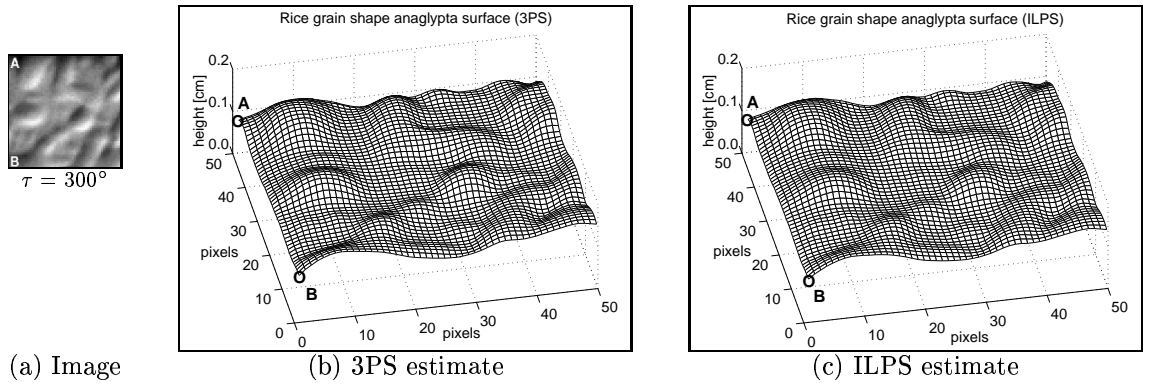


Figure C.34: Height recovery for the *rice grain shape anaglypta*

# References

- [1] D-M Tsai and B. Hsiao. Automatic surface inspection using wavelet reconstruction. *Pattern Recognition*, 34(6):1285–1305, June 2001.
- [2] T. Fricke-Begemann, G. Gulker, K.D. Hinsch, and K. Wolff. Corrosion monitoring with speckle correlation. *Applied Optics*, 38(28):5948–5955, October 1999.
- [3] C.E. Costa and M. Petrou. Automatic registration of ceramic tiles for the purpose of fault detection. *Machine Vision and Applications*, 11(5):225–230, 2000.
- [4] Y-J Chang and Y-C Chen. Facial model adaptation from a monocular image sequence using a textured polygonal model. *Signal Processing: Image Communication*, 17(5):373–392, May 2002.
- [5] M. Bresch. Optimizing filter banks for supervised texture recognition. *Pattern Recognition*, 35(4):783–790, April 2002.
- [6] M.J. Chantler and G. McGunnigle. On the use of gradient space eigenvalues for rotation invariant texture classification. *Proceedings of the 15th International Conference on Pattern Recognition, ICPR-2000*, 3:931–934, September 2000.
- [7] J.W. Hermann, J. Radice, J. Disciullo, and P. Chipman. Flexible lighting modules machine vision. *Advance Packaging*, pages 27–32, May 2000.
- [8] E.N. Coleman and R. Jain. Obtaining 3-dimensional shape of textured and specular surfaces using a four-source photometry. *Computer Graphics and Image Processing*, 18:309–328, 1982.
- [9] G.C. Bakos, N.F. Tsagas, J. Lygouras, and J. Lucas. Long-distance noncontact high-precision measurements. *International Journal of electronics*, 75(6):1269–1279, December 1993.
- [10] J.S. Massa, G.S. Buller, A.C. Walker, S. Cova, M. Umasuthan, and A.M. Wallace. Time-of-flight optical ranging system based on time-correlated single-photon counting. *Applied Optics*, 37(31):7298–7304, November 1998.
- [11] K.E. Torrance and E.M. Sparrow. Theory for off-specular reflection from roughened surfaces. *Optical Society of America*, 57(9):1105–14, 1967.
- [12] O. Fauberas, editor. *Three-dimensional computer vision: a geometric viewpoint*. Cambridge MA.: M.I.T. Press, 1995.
- [13] B. Shahraray and M.K. Brown. Robust depth estimation from optical flow. *Second International Conference on Computer Vision, IEEE Comput. Soc. Press*, pages 641–650, December 1988.
- [14] D.A. Forsyth. Shape from texture and integrability. *Proceedings of the 8th IEEE International Conference On Computer Vision*, 2:447–452, July 2001.
- [15] M. Proesmans, L. Van Gool, and F. Defoort. Reading between the lines - a

- method for extracting dynamic 3d from texture. *Proceedings of the 6th IEEE International Conference On Computer Vision*, 1:1081–1086, January 1998.
- [16] E.P. Krotkov. Focusing. *International Journal of Computer Vision*, 1(3):223–237, October 1987.
  - [17] E.P. Krotkov and R. Bajcsy. Active vision for reliable ranging: cooperating focus, stereo and vergence. *International Journal of Computer Vision*, 11(2):187–203, October 1993.
  - [18] B.K.P. Horn, editor. *Robot Vision*. The MIT Press, 1 edition, 1986.
  - [19] D. Nandy and J. Ben-Arie. Shape from shading: a novel approach for 3-d face shape recovery. *IEEE Transactions on Image Processing*, 10(2):206–217, February 2001.
  - [20] J.H. Lambert, editor. *Photometria*. Augsburg, 1760.
  - [21] J.M. Bennett and L. Mattson, editors. *Introduction to surface roughness and scattering*. Optical Society of America, 2 edition, 1999.
  - [22] E. Kreyszig, editor. *Advanced engineering mathematics*. John Wiley and Sons, 5 edition, 1983.
  - [23] J.A. Ogilvy, editor. *Theory of wave scattering from random rough surfaces*. Adam Hilger, 1991.
  - [24] A.P. Pentland. Shading into texture. *Artificial Intelligence*, 29(2):147–170, August 1986.
  - [25] D.C. Knill and D. Kersten. Learning a near-optimal estimator for surface shape from shading. *Computer Vision, Graphics, and Image Understanding*, 50(1):75–100, April 1990.
  - [26] A. Conci and C.B. Proenca. A fractal image analysis system for fabric inspection based on a box-counting method. *Computer Networks and Industrial Systems*, 30(12):1887–95, November 1998.
  - [27] F. Espinal, T. Huntberger, B.D. Jawerth, and T. Kubota. Wavelet-based fractal signature analysis for automatic target recognition. *Optical Engineering*, 37(1):166–174, January 1998.
  - [28] B.B. Mandelbrot, editor. *Fractals: Form, Chance, and Dimension*. Freeman, 1977.
  - [29] A.P. Pentland. Fractal-based description of natural scenes. *IEEE Transaction on Pattern Analysis and Machine Intelligence, PAMI*, 6:661–674, November 1984.
  - [30] A. Bunde and S. Havlin, editors. *Fractals and Disordered Systems*. Springer-Verlag, 1991.
  - [31] D.J. Mulvaney, D.E. Newland, and K.F. Gill. A complete description of surface texture profiles. *Wear*, 132:173–182, 1989.
  - [32] M.S. Keshner. 1/f noise. *Proceedings of the IEEE*, 70(3):212–218, March 1982.
  - [33] J.J. van Wijk. Spot noise. *Computer Graphics*, 25(4):309–318, July 1991.
  - [34] A. Watt, editor. *3D computer graphics*. Addison-Wesley, 3 edition, 1993.
  - [35] J.J. Koenderink and A.J. van Doorn. Bidirectional reflection distribution function expressed in terms of surface scattering modes. *Proceeding of the 4th European Conference on Computer Vision*, 2:28–39, 1996.
  - [36] K.J. Dana, B. van Ginneken, S.K. Nayar, and J.J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, 1999.
  - [37] J.J. Koenderink and A.J. van Doorn. Illuminance texture due to surface

- mesostructure. *Journal of the Optical Society of America A*, 13:452–463, 1996.
- [38] S. Magda and D.J. Kriegman. Beyond lambert: reconstructing surfaces with arbitrary brdfs. *Eighth International Conference on Computer Vision*, 2:391–398, 2001.
  - [39] L.B. Wolff. Generalizing lambert’s law for smooth surfaces. *Proceedings of Fourth European Conference on Computer Vision*, 2:40–53, April 1996.
  - [40] M. Oren and S.K. Nayar. Generalization of the lambertian model and implication for machine vision. *International Journal of Computer Vision*, 14:227–251, 1995.
  - [41] T-Y Chao, S-J Wang, and H-M Hang. Stereoscopic correspondence by applying physical constraints and statistical observations to dissimilarity map. *Proceeding of Spie - The International Society for Optical Engineering*, 3957:78–89, January 2000.
  - [42] U.R. Dhond and J.K. Aggarwal. A cost-benefit analysis of a third camera for stereo correspondence. *International Journal on Computer Vision*, 6(1):39–58, 1991.
  - [43] Sandmeier S.R. Acquisition of bidirectional reflectance factor data with field goniometers. *Remote Sensing of Environment*, 73(3):257–269, September 2000.
  - [44] B.T. Phong. Illumination for computed generated pictures. *Communications of the ACM*, 18(6):311–317, June 1975.
  - [45] G Healey and T.O. Binford. Local shape from specularity. *Proceedings on the First IEEE International Conference on Computer Vision*, 1:151–60, 1987.
  - [46] K. Ikeuchi and K. Sato. Determining reflectance parameters using range and brightness images. *Proceedings on the Third IEEE International Conference on Computer Vision*, 1:12–20, 1990.
  - [47] F. Solomon and K. Ikeuchi. 4-source method for determining reflectance roughness. *Proceedings on the IEEE Conference on Computer Vision and Pattern Recognition*, 1:466–71, 1992.
  - [48] F. Solomon and K. Ikeuchi. Extracting the shape and roughness of specular lobe objects using four light photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(4):449–54, 1996.
  - [49] L.B. Wolff and T.E. Boult. Constraining object features using polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):635–57, 1991.
  - [50] J.F. Blinn. Models of light reflection for computer synthesized pictures. *Computer Graphics*, 11(2):192–8, 1977.
  - [51] R.L. Cook and K.E. Torrance. A reflectance model for computer graphics. *Computer Graphics*, 15(3):307–16, 1982.
  - [52] P. Beckmann and A. Spizzichino, editors. *Scattering of electromagnetic waves from rough surfaces*. New York: Pergamon, 1963.
  - [53] S.K. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):611–634, July 1991.
  - [54] A.J. Lundberg, L.B. Wolff, and D.A. Socolinsky. New perspectives on geometric reflection theory from rough surfaces. *Proceedings on the Eighth IEEE International Conference on Computer Vision*, 1(1):225–32, 2001.
  - [55] P. Kube and A. Pentland. On the imaging of fractal surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(3):704–707, September

1988.

- [56] A.P. Pentland. Linear shape from shading. *International Journal of Computer Vision*, 4:153–162, 1990.
- [57] M.J. Chantler. *The effect of variation in illuminant direction on texture classification*. PhD thesis, PhD Thesis, Department of Computing and Electrical Engineering, Heriot-Watt University, 1994.
- [58] G. McGunnigle and M.J. Chantler. Evaluating kube and pentland’s fractal imaging model. *IEEE Transactions on Image Processing*, 10(4):534–42, April 2001.
- [59] G. McGunnigle. *The classification of textured surfaces under varying illuminant direction*. PhD thesis, PhD Thesis, Department of Computing and Electrical Engineering, Heriot-Watt University, 1998.
- [60] P Dupuis and J .Oliensis. Direct methods for reconstructing shape from shading. *IEEE Proceeding on Computer Vision and Pattern Recognition*, pages 453–458, 1992.
- [61] J Oliensis and P. Dupuis. A global algorithm for shape from shading. *Proceeding on the International Conference on Computer Vision*, pages 692–701, 1993.
- [62] M. Bichsel and A.P. Pentland. A simple algorithm for shape from shading. *IEEE Proceeding on Computer Vision and Pattern Recognition*, pages 459–465, 1992.
- [63] R. Kimmel and A.M.. Bruchstein. Traching level sets by level sets: a method for solving the shape from shading problem. *Computer Vision and Image Understanding*, 62(1):47–58, July 1995.
- [64] K. Ikeuchi and B.K.P. Horn. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17(1-3):141–184, 1981.
- [65] M.J. Brooks and B.K.P. Horn. Shape and source from shading. *Proceedings International Joint Conference Artificial Intelligence*, pages 932–936, 1985.
- [66] 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):439–451, July 1988.
- [67] B.K. Horn. Height and gradient from shading. *International Journal of Computer Vision*, 5(1):37–75, 1990.
- [68] R. Szeliski. Fast shape from shading. *Computer Vision, Graphics, Image Processing: Image Understanding*, 53:139–153, 1991.
- [69] Q. Zheng and R. Chellappa. Estimation of illuminant direction, albedo and shape from shading. *IEEE Transaccions Pattern Analysis and Machine Inteligent*, 13(7):680–702, 1991.
- [70] R. Zhang and M. Shah. Shape from intensity gradient. *IEEE Transactions on System, Man and Cybernetics Part A: Systems and Humans*, 29(3):318–325, May 1999.
- [71] Y.G. Leclerc and A. F. Bobick. The direct computation of height from shading. *IEEE Proceeding on Computer Vision and Pattern Recognition*, pages 552–558, 1991.
- [72] K.M. Lee and C.C.J. Kou. Shape from shading with a linear triangular surface model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(8):815–822, August 1993.
- [73] W Guo-Qing and G Hirzinger. Parametric shape-from-shading by radial basis

- functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(4):353–365, April 1997.
- [74] P.L. Worthington and E.R. Hancock. Surface topography using shape-from-shading. *Pattern Recognition*, 34(4):823–840, April 2001.
  - [75] S-Y Cho and T.W.S. Chow. Shape and measurement technology by an improved shape-from-shading neural algorithm. *IEEE Transactions on Industrial Electronics*, 47(1):225–230, February 2000.
  - [76] P.S. Tsai and M. Shah. Shape from shading using linear approximation. *Image and Vision Computing Journal*, 12(8):487–498, 1994.
  - [77] P.S. Tsai and M. Shah. Shape from shading with variable albedo. *Optical Engineering*, 37(4):1212–1220, April 1998.
  - [78] R. Kozera and R. Klette. Well-posedness of linear shape-from-shading problem. *7th International Conference on Computer Analysis of Images and Patterns.*, pages 10–12, September 1997.
  - [79] B.K.P. Horn and M.J. Brooks, editors. *Shape from Shading*. The MIT Press, 1 edition, 1989.
  - [80] M.J. Brooks, W. Chojnaki, and R. Kozera. Circularly-symmetric eikonal equations and non-uniqueness in computer vision. *Journal of Mathematical Analysis and Applications*, 165(1):192–215, March 1982.
  - [81] J Oliensis. Uniqueness in shape from shading. *International Journal of Computer Vision*, 6(2):74–104, 1991.
  - [82] P. Maragos. Differential morphology and image processing. *IEEE Transactions of Image Processing*, 5(6):922–937, 1996.
  - [83] R. Kimmel and J.A. Sethian. Optimal algorithm for shape from shading and path planning. *Journal of Mathematical Imaging and Vision*, 14(3):237–244, May 2001.
  - [84] R.J. Woodham. Photometric stereo: a reflectance map technique for determining surface orientation from image intensity. *Proceeding of Society of Photo-Optical Instrumentation Engineers*, 155:136–143, 1978.
  - [85] T. Kulick. Shape from shading using three images. *Computing*, 57(1):1–24, 1996.
  - [86] B-H Kim and R-H Park. Shape from shading and photometric stereo using surface approximation by legendre polynomials. *Computer Vision and Image Understanding*, 66(3):255–270, June 1997.
  - [87] R. Zhang and M. Shah. Iterative shape recovery from multiple images. *Image and Vision Computing*, 15(11):801–814, November 1997.
  - [88] G. McGunnigle and M.J. Chantler. Segmentation of rough surfaces using reflectance. *British Machine Vision Conference, BMVC*, 1:323–332, September 2001.
  - [89] K. Ikeuchi. Determining surface orientation of specular surfaces by using the photometric stereo method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-3(6):661–669, 1981.
  - [90] S.K. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric stereo. *IEEE Transactions on Robotics and Automation*, 6(4):418–431, August 1990.
  - [91] H.D. Tagare and R.J.P. deFigueiredo. A theory of photometric stereo for a class of diffuse non-lambertian surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(2):133–152, February 1991.

- [92] G. Kay and T. Caelli. Estimating the parameters of an illumination model using photometric stereo. *Graphical Models and Image Processing*, 57(5):365–388, September 1995.
- [93] R.J. Woodham. Gradient and curvature from the photometric stereo method, including local confidence estimation. *Journal of the Optical Society of America, A Optical Image, Science and Vision*, 11(11):3050–3068, November 1994.
- [94] Y. Iwahori, R.J. Woodham, M.S. Bhuiyan, and N. Ishii. Neural network based photometric stereo for object with non-uniform reflectance factor. *Proceeding of the IEEE International Conference on Neural Information Processing*, 3:1213–1218, 1999.
- [95] K. Schluns and O. Wittig. Photometric stereo for non-lambertian surfaces using colour information. *Proceedings on the 5th International Conference on Computer Analysis of Images and Patterns*, 1:444–451, September 1993.
- [96] H. Christensen and L.G. Shapiro. Three-dimensional shape from colour photometric stereo. *International Journal on Computer Vision*, 13(2):213–227, 1994.
- [97] A.A. Mustafa, L.G. Shapiro, and M.A. Ganter. 3d object recognition from colour intensity images. *IEEE Proceedings of International Conference on Pattern Recognition*, 1:627–631, August 1996.
- [98] A.P. Petrov and K. Antonova. Resolving the colour image irradiance equation. *Colour, Research and Applications*, 21(2):97–103, April 1996.
- [99] M.S. Drew and M.H. Brill. Colour from shape from colour: a simple formalism with known light sources. *Journal of Optical Society of America A-Optics and Image Science*, 17(8):1371–1381, August 2000.
- [100] S. Barsky and M. Petrou. Colour photometric stereo: simultaneously reconstruction of local gradient and colour of rough textured surfaces. *Proceedings of the eighth IEEE International Conference on Computer Vision, ICCV*, 2:600–605, 1002.
- [101] J.J. Clark. Active photometric stereo. *IEEE Proceeding of the Computer Vision and Pattern Recognition Conference*, 1:29–34, June 1992.
- [102] Y. Iwahori, R.J. Woodham, and N. Ishii. Shape from shading with a nearby moving point light source. *Proceeding on the second International Conference on Automation, Robotics and Computer Vision*, 5(5):1–5, 1992.
- [103] Y. Iwahori, R.J. Woodham, H. Tanaka, and N. Ishii. Moving point light source photometric stereo. *IEICE Transactions on Information and Systems*, E77-D(8):925–929, 1994.
- [104] J. J. Clark and H. Pekau. An integral formulation for differential photometric stereo. *IEEE Proceeding of the Computer Vision and Pattern Recognition Conference*, 1:129–134, 1999.
- [105] T. Okatani and K. Deguchi. On uniqueness of solutions of the three-light-source photometric stereo: Conditions on illumination configuration and surface reflectance. *Computer Vision and Image Understanding*, 81(2):211–226, February 2001.
- [106] R. Klette and K. Schluns. Height data from gradient fields. *Proceedings of Machine Vision: Applications, Architectures and Systems Integration V, SPIE*, 2908:204–215, 1996.
- [107] G. Healey and R. Jain. Depth recovery from surface normals. *Proceeding 7th International Conference on Pattern Recognition, IEEE*, 2:894–896, 1984.

- [108] Z. Wu and L. Li. A line-integration based method for depth recovery from surface normals. *Computer Vision, Graphics and Image Processing*, 43:53–66, 1988.
- [109] K. Ikeuchi. Reconstructing a depth map from intensity images. *IEEE 7th International Conference on Pattern Recognition*, 2:736–738, 1984.
- [110] T. Simchony, R. Chellappa, and M. Shao. Direct analytical methods for solving poisson equations in computer vision problems. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 12(5):435–446, May 1990.
- [111] N. Petrovic, I. Cohen, B.J. Frey, R. Koetter, and T.S. Huang. Enforcing integrability for surface reconstruction algorithms using belief propagation in graphical models. *Proceeding of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2001*, 1:743–748, 2001.
- [112] J.G. Harris. A new approach to surface reconstruction: the coupled depth/slope model. *Proceedings of the 1st IEEE Conference on Computer Vision*, 1:277–283, June 1987.
- [113] G. McGunnigle and M.J. Chantler. Rotation insensitive classification of rough surfaces. *IEE Proceedings Vis. Image and Sig*, 146(6):345–352, December 1999.
- [114] G. McGunnigle and M.J. Chantler. Rough surfaces classification using point statistics from photometric stereo. *Pattern Recognition Letters*, 21(6-7):593–604, 2000.
- [115] R. Zhang, P.S. Tsai, J.E. Cryer, and M. Shah. Shape from shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706, August 1999.
- [116] P. Hansson and P.A. Johansson. Topography and reflectance analysis of paper surfaces using photometric stereo method. *Optical Engineering*, 39(9):2555–2561, September 2000.
- [117] G. McGunnigle and M.J. Chantler. Recovery of fingerprints using photometric stereo. *Irish Machine Vision and Image Processing Conference*, pages 192–199, September 2001.
- [118] G. McGunnigle and M.J. Chantler. *Recovery of indented handwriting*. PhD thesis, Internal Research Memorandum, RM/02/3, Department of Computing and Electrical Engineering, Heriot-Watt University, 2002.
- [119] G. McGunnigle and M.J. Chantler. *Photometric Stereo and oilpainting: techniques and applications*. PhD thesis, Internal Research Memorandum, RM/02/1, Department of Computing and Electrical Engineering, Heriot-Watt University, 2002.
- [120] R. Onn and A. Bruckstein. Integrability disambiguities surface recovery in two-image photometric stereo. *International Journal of Computer Vision*, 5(1):105–113, 1990.
- [121] R. Kozera. On shape recovery from two shading patterns. *International Journal of Pattern Recognition and Artificial Intelligence*, 6(4):673–698, 1992.
- [122] K.M. Lee and C.C.J. Kou. Surface reconstruction from photometric stereo images. *Journal of the Optical Society of America A-Optics and Image Science*, 10(5):855–868, May 1993.
- [123] J.R.A. Torreão and J.L. Fernandes. Matching photometric stereo images. *Optical Society of America*, 15(2):2966–2975, December 1998.
- [124] J.R.A. Torreão. A green’s function approach to shape from shading. *Pattern Recognition*, 34(12):2367–2382, December 2001.



- [125] J. Yang, N. Ohnishi, and N. Sugie. Two image photometric stereo method. *Proceedings of Spie - The International Society for Optical Engineering*, 1826:452–463, November 1992.
- [126] A.E. Caver and R.J. Schalkoff. Deformation detection on composite reflectance surfaces using two image photometric stereo. *Optical Engineering - Society of Photo-Optical Instrumentation Engineers*, 40(2):295–302, February 2001.
- [127] B-Y Kim and R-H Park. Multi-image photometric stereo using surface approximation by legendre polynomials. *Pattern Recognition*, 31(8):1033–1047, May 1998.
- [128] D.H. Ballard and C.M. Brown, editors. *Computer Vision*. Prentice-Hall, 1982.
- [129] M.J. Bentum, B.A. Barthoold, and T. Malzbender. Frequency analysis of gradient estimators in volume rendering. *IEEE Transaction on Visualization and Computer Graphics*, 2(3):242–254, September 1996.
- [130] R. Porter and N. Canagarajah. Robust rotation-invariant texture classification: wavelet, gabor filter and gmrf based schemes. *IEE Proceedings on Vision, Image and Signal Processing*, 144(3):180–188, June 1997.
- [131] S.F. Gilyazov and N.L. Gol’dman, editors. *Regularization of ill-posed problems by iteration methods*. Klumer Academic Publisher, 2000.
- [132] H. Hayakawa, Y. Wada, and M. Kawato. Solution of nonlinear vision problem based on forward and approximated inverse optics models. *Electronics and Communications in Japan*, 78(5):22–33, 1995.
- [133] M. Subbarao and Y.F. Liu. Analysis of defocused image data for 3d shape recovery using a regularization technique. *The International Society of Optical Engineering - SPIE*, 3204:24–35, 1997.
- [134] F.P. Ferrie and J. Lagarde. Curvature consistency improves local shading analysis. *Proceedings of the IEEE International Conference on Pattern Recognition*, I:70–76, 1990.
- [135] P.L. Worthington and E.R. Hancock. New constraints on data-closeness and needle map consistency for shape-from-shading. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 21(12):1250–1267, December 1999.
- [136] Smith W.A.P. and E.R. Hancock. Face recognition using shape-from-shading. *British Machine Vision Conference*, 2:597–606, September 2002.