Improved Shape from Shading using Non-Lambertian Reflectance Models

Hossein Ragheb

Submitted for the degree of Philosophy Doctor

Department of Computer Science

THE UNIVERSITY of York

May 17, 2004

Abstract

This thesis focuses on the use of non-Lambertian reflectance models for curved surface analysis. The contributions are: a) to show how specular and diffuse reflectance models can be used to recover surface normal estimates using shape-from-shading (SFS), b) to show how wave scattering theory can be used to estimate rough surface characteristics, and c) to present a method for surface normal smoothing that imposes consistency on principal curvature directions.

One of the problems that hinders standard methods for SFS is the presence of local specularities which may be misidentified as high curvature surface features. To address this problem we develop a maximum *a posteriori* probability estimation method for estimating the mixing proportions for Lambertian and specular reflectance, and also, for recovering local surface normals. The framework for our study is provided by the iterated conditional modes algorithm. Once improved estimates of surface normal directions are obtained, we can also remove the specular highlights from the raw image using the reconstructed specular intensity image.

It is also well known that the non-specular reflectance from both rough and shiny surfaces depart significantly from the predictions of Lambert's law. These effects are particularly marked at the occluding boundary where the surface is highly inclined to the viewer direction. Hence, Lambertian SFS methods can not be applied directly to the analysis of rough and shiny surfaces. In order to overcome this difficulty, we consider how to reconstruct the Lambertian component when the object is illuminated in the viewing direction. To do this we make use of the diffuse reflectance models described by Oren and Nayar, and by Wolff.

Next, the thesis focuses on how reflectance models based on wave scattering theory and reported in the physics literature can be used for rough surface analysis. For dielectrics, the models investigated are derived from the Beckmann-Kirchhoff scatter theory, while for metals the starting point is the Davies scattering model. We show how these models can be used to estimate surface roughness parameters using simple reflectance

measurements. With estimates of these parameters to hand, we take our analysis one step further by demonstrating how the Vernold-Harvey modification of the Beckmann model may be used for curved surface analysis.

Finally, the thesis describes a new surface normal smoothing process which can be used in conjunction with SFS. Rather than directly smoothing the surface normal vectors, we exert control over their directions by smoothing the field of principal curvature vectors. To do this we develop a topography sensitive smoothing process which overcomes the problems of singularities in the field of principal curvature directions at the locations of umbilics and saddles. We also investigate the effect of using the Woodham's viewpoint-invariant estimator of the Hessian in the smoothing process and in the final surface normal directions.

Contents

1 Introduction **2**

2 Literature Review **7**

2.1 Shape from Shading 8

2.1.1 Curvature Consistency 9

2.2 Reflectance Modelling 10

2.2.1 What Roughness Means 11

2.2.2 Specular Reflectance Models 13

2.2.3 Physics-Based Reflectance Models 15

2.2.4 Phenomenological Reflectance Models 19

2.3 Other Related Topics 20

2.3.1 Multiple Light Sources 21

2.3.2 Shading and Illumination 22

2.3.3 Textured Surfaces 23

2.4 Conclusions 25

3 Highlight Removal using Specular Shape-from-Shading **27**

3.1 Reflectance Geometries and Models 29

3.1.1 Specular Reflectance 29

3.1.2 Lambertian Reflectance 33

3.2 Bayesian Framework 34

3.2.1 Formal Details 35

3.2.2 Algorithmic Steps 39
3.3 Distribution of Specular Component 40
3.3.1 Beckmann Distribution 41
3.3.2 Gaussian Distribution 43
3.3.3 Intensity Distribution 43
3.4 Distribution of Lambertian Component 44
3.5 Smoothness Prior 44
3.6 Specularity Removal 45
3.7 Experiments 47
3.7.1 Examples 47
3.7.2 Comparison 52
3.7.3 Specular Detail 54
3.7.4 A Posteriori Probabilities 54
3.7.5 Variable Light-Source Direction 63
3.7.6 Ground-Truth Data 64
3.8 Conclusions 68

4 Non-Lambertian Diffuse Reflectance Models for Photometric Correction 69

4.1 Diffuse Reflectance Models 71
4.1.1 Oren-Nayar Model for Rough Surfaces 72
4.1.2 Wolff Model for Smooth Surfaces 77
4.1.3 Combined Model for Intermediate Roughness 79
4.2 Photometric Correction Process 80
4.2.1 Analytic Solution 81
4.2.2 Lookup Table Solution 82
4.3 Experiments 83
4.3.1 Synthetic Data 84
4.3.2 Real-world Data 90
4.4 Conclusions 103

5 Physics-Based Reflectance Models for Rough Surface Analysis 104

5.1 Beckmann-Kirchhoff Scatter Theory 106

5.1.1 Modified Beckmann-Kirchhoff Model 109

5.1.2 Domain of Validity 112

5.1.3 Comparison with Phenomenological Reflectance Models 113

5.2 Davies Scatter Theory 116

5.2.1 Bennett-Porteus Formulation 116

5.2.2 Torrance Formulation 118

5.3 Surface Characteristic Estimation 119

5.3.1 Dielectric Surfaces with Small Roughness 121

5.3.2 Dielectric Surfaces with Large Roughness 124

5.3.3 Metallic Surfaces with Small Roughness 125

5.4 Exploiting the Beckmann Model 126

5.4.1 Photometric Correction 126

5.4.2 Surface Normal Recovery and Re-illumination 130

5.5 Experiments 130

5.5.1 Rough Surface Parameter Estimation 131

5.5.2 Comparing the Beckmann model with Reflectance Data 135

5.5.3 Photometric Correction 137

5.5.4 Shape-from-Shading and Re-illumination 140

5.5.5 Physics-Based Rendering 145

5.6 Conclusions 150

6 Curvature Consistency using Darboux Smoothing 151

6.1 Local Darboux Frames 153

6.2 Viewpoint-Invariant Curvatures 155

6.3 Smoothing Principal Curvature Directions 156

6.4 Surface Normal Adjustment 157

6.4.1 Rotation About the Irradiance Cone 158

6.4.2 Computing the Darboux Surface Normal Direction 159
6.4.3 Imposing the Irradiance Constraint 160
6.5 Experiments 161
6.5.1 Synthetic Data 162
6.5.2 Real-World Data 168
6.5.3 Hessian Compensation 170
6.6 Conclusions 177

7 Conclusions **178**

7.1 Strengths and Weaknesses of the Thesis 178
7.2 Suggestions for Further Research 181
7.2.1 Shape from Multiple Images using Physical Reflectance Models . 182
7.2.2 Roughness Estimation for the Oren-Nayar Model 184
7.2.3 Testing Alternative Reflectance Models 185
7.2.4 Modifying the Beckmann Model for Missing Effects 186
7.2.5 Reflectance Data Against Reflectance Models 187

List of Tables

5.1 Results for the slightly-rough metallic surfaces shown in Figure 5.13. . . 133

5.2 Results for the slightly-rough dielectric surfaces shown in Figure 5.11. . . 133

5.3 Results for the very-rough dielectric surfaces shown in Figure 5.14. . . . 135

6.1 Topographic classes and the corresponding shape-index intervals. 155

List of Figures

3.1 A simple geometry of the specular reflection process. 30

3.2 Geometry of the surface normal updating process: a) initialization of Lambertian surface normal $\vec{N}_L^{(0)}$ (initial posterior-mean surface normal $\vec{M}^{(0)}$ is identical to $\vec{N}_L^{(0)}$); b) smoothing the current posterior-mean surface normal $\vec{M}^{(n)}$ to $\vec{M}_R^{(n)}$, and, computing $\vec{N}_L^{(n)}$ by rotating $\vec{M}_R^{(n)}$ onto the cone ($\vec{N}_S^{(n)}$ is fixed to $\vec{N}_S$); c) computing the updated posterior-mean surface normal $\vec{M}^{(n+1)}$ using the current Lambertian and specular surface normals ($\vec{N}_L^{(n)}$ and $\vec{N}_S^{(n)}$) and the corresponding *a posteriori* probabilities. 39

3.3 Dependence of the Beckmann distribution on the parameter σ_m. 42

3.4 Highlight removal and shape recovery for a porcelain bear. 49

3.5 Highlight removal and shape recovery for a porcelain vase. 50

3.6 Highlight removal and shape recovery for a porcelain urn. 51

3.7 Comparing results corresponding to the raw and separated matte images. . 52

3.8 Comparing results corresponding to the raw and separated matte images. . 53

3.9 Comparing results corresponding to the raw and separated matte images. . 53

3.10 The porcelain urn. 55

3.11 Central highlight on the porcelain urn. 55

3.12 Specularities on the handle of the porcelain urn. 55

3.13 Re-illuminated images of the handle of the porcelain urn for incidence angles from -40° to $+40^\circ$. 55

3.14 Results for the central detail on the urn (Figure 3.11) using different distribution functions: (a) the raw and matte grey-levels, (b) specular probability scatter plots. 56
3.15 Histograms of the (*a posteriori*) specular probabilities for the porcelain vase using different distribution functions and parameter values (σ_L is fixed at 0.1). 57
3.16 Histograms of the (*a posteriori*) specular probabilities for the porcelain bear using different distribution functions and parameter values (σ_L is fixed at 0.1). 58
3.17 Results for highlights on the neck of the vase using different distribution functions (specular patches are enlarged to visualize grey-scale levels of the pixels). 60
3.18 Results for highlights on the bulb of the vase using different distribution functions (specular patches are enlarged to visualize grey-scale levels of the pixels). 61
3.19 Results for highlights on the left leg of the bear using different distribution functions (specular patches are enlarged to visualize grey-scale levels of the pixels). 62
3.20 Removing specular highlights from images of the vase taken under different light-source directions, while the viewing direction is aligned with the z axis. 64
3.21 Pairwise difference images before specularity removal. 65
3.22 Pairwise difference images after specularity removal. 65
3.23 Residual differences between difference images of Figures 3.21 and 3.22. 66
3.24 Needle maps of (1) recovered surface normals for separated matte images of Figure 3.20; (2) corresponding fields of difference vectors. 66
3.25 Surface normals for the raw and matte images compared to the ground-truth. 67

4.1 Typical geometry for the local tangent plane P_N: the normal $\vec{N}$ to the surface, the light-source direction $\vec{L}$, the viewing direction $\vec{V}$ and the corresponding angles. 74

4.2 Reflectance geometries used to extract simplified equations from the Oren-Nayar model for the possible special cases; left to right: (a-d) Cases 1-4. . 75

4.3 Surface radiance versus incidence angle using the simplified model of Oren-Nayar (Equation 4.9) for $\sigma_a = 0.05, 0.2, 0.4, 0.6, 0.8, 1.0$rd. 77

4.4 Surface radiance versus incidence angle using the simplified model of Wolff (Equation 4.13) for $n = 1.4, 1.5, 1.7, 1.9, 2.0$. 79

4.5 Surface radiance versus incidence angle using the simplified combined model (Equation 4.15) for $n = 1.7$ and $\sigma_a = 0.01, 0.1, 0.15, 0.2, 0.25, 0.3$rd. 81

4.6 Synthetic surface: (a) elevation data (b) field of surface normal directions. 84

4.7 Images of the synthetic surface of Figure 4.6 rendered using the models of (a) Lambert, (b) Wolff ($n = 1.7$), (c) Oren-Nayar ($\sigma_a = 20^\circ$), and (d) Wolff-Nayar-Oren ($n = 1.7$ and $\sigma_a = 20^\circ$). 84

4.8 Applying Lambertian SFS (Chapter 6) to images of Figure 4.7: (1) recovered surface normals; (2) difference between the recovered and ground-truth surface normals. 85

4.9 (1) Zenith angles and (2) azimuth angles of the recovered surface normals versus the corresponding ground-truth angles. 86

4.10 TMS surface normal error versus iteration number obtained by applying Lambertian SFS (Chapter 6) to the images of Figure 4.7 generated using different models. 86

4.11 (1) The corrected images obtained by applying the table lookup method to the images of Figure 4.7 using the (a) Wolff, (b) Oren-Nayar, and (c) combined models; (2) difference images. 87

4.12 (1) Images obtained by adding Gaussian noise to images of Figure 4.7; (2) corrected images obtained by applying the table lookup method to the noisy images using the models of (a) Wolff, (b) Oren-Nayar, and (c) Wolff-Nayar-Oren; (3) difference images. 88
4.13 Images of arc-cosine (ACS) error between the ground-truth and the estimated surface normals corresponding to the (1) uncorrected images of Figure 4.7, (2) corrected images of Figure 4.11, (3) uncorrected and (4) corrected noisy images of Figure 4.12. 89
4.14 TMS surface normal error versus iteration number obtained by applying Lambertian SFS (Chapter 6) to the uncorrected (dashed) and corrected (solid) images shown in (1) Figures 4.7 and 4.8, and, (2) in Figure 4.12. . 90
4.15 Removing (b) the specular component from (a) the raw image to obtain (c) the diffuse component using the method described in Chapter 3. 91
4.16 Photometric correction for (a) the porcelain vase using the Wolff model, (b) a terra-cotta bear using the combined model, and (c) a chalk owl using the Oren-Nayar model: (1) diffuse images; (2) corrected Lambertian images; (3) difference images; (4) transverse brightness sections. 93
4.17 The surface normals obtained by applying Lambertian SFS (1) to the diffuse images of Figure 4.16, and (2) to their corresponding corrected Lambertian images; (3) the field of difference. 94
4.18 Scatter plots for the incidence (a) zenith angles and (b) azimuth angles (degrees) extracted from the recovered surface normals for the terra-cotta bear: the horizontal axis shows the angles corresponding to the diffuse image, while the vertical axis shows those for the Lambertian image (using the Oren-Nayar model); (c) the difference between the diffuse and Lambertian azimuth angles versus the Lambertian zenith angle. 96
4.19 Scatter plots as described in Figure 4.18 for the porcelain vase using the Wolff model. 97

4.20 (1) Strawboard cylinder: (a) raw image; corrected images using the Oren-Nayar model (b) for $\sigma_a = 0.1$rd and (c) for $\sigma_a = 0.12$rd; (2) magnified difference between the raw and corrected images. 97
4.21 (a) Lambertian image for a fitted synthetic cylinder, and its absolute differences with (b) the raw and (c-d) corrected images of Figure 4.20. . . . 97
4.22 (a) Normalized raw radiance for the strawboard cylinder versus θ_i compared to the corrected radiance using the Oren-Nayar model ($\sigma_a = 0.1$ and $\sigma_a = 0.12$rd) and the Lambert's law; (b) TMS surface normal error versus iteration number obtained by applying SFS to the images of Figure 4.20. 98
4.23 Removing specular highlights (using the method described in Chapter 3) from the raw image for the neck of a porcelain vase. 100
4.24 (1) Cylindrical section of the vase: (a) non-specular image, corrected images using the (b) Wolff ($n = 1.7$), (c) Oren-Nayar ($\sigma_a = 0.1$rd), and (d) combined models; (2) magnified differences between the non-specular and corrected images. 100
4.25 (a) Lambertian image for a fitted synthetic cylinder and its absolute differences with the (b) non-specular and (c) corrected images of Figure 4.24. 100
4.26 TMS surface normal error versus iteration number obtained by applying SFS to the non-specular and corrected images of the neck of the vase (Figure 4.24). 101
4.27 Re-illumination of the recovered surface normals (Figure 4.17) to generate diffuse images for the porcelain vase using (a) the Wolff model, (b) the combined model and (c) the Oren-Nayar model for three light-source directions: (1) $(0, 0)$, (2) $(45°, -90°)$, and (3) $(90°, -90°)$; the parameters of the models are set to $\sigma_a = 0.15$rd and $n = 2.0$. 102

5.1 Scattering geometry on the tangent plane coordinate system: the incident light propagates in the direction $\vec{L}$, while the direction for the scattered light is $\vec{V}$. 107

5.2 Normalized intensity using the original B-K model with (a) a Gaussian or (b) an exponential correlation function versus θ_i for several values of $\eta = (T/\sigma)^2$; Lambert's law (dashed). 111

5.3 Plots (as described in Figure 5.2) for the modified B-K model. 111

5.4 Comparing the modified B-K model (Equation 5.9) with the Oren-Nayar model (Equation 5.10): (a) $\sigma/T = 1.5$, $\sigma_a = 0.3$rd; (b) $\sigma/T = 1.8$, $\sigma_a = 0.6$rd. 114

5.5 Comparing the modified B-K model (Equation 5.9) with the combined model (Equation 5.11, $n = 1.7$): (a) $\sigma/T = 1.4$, $\sigma_a = 0.2$rd; (b) $\sigma/T = 1.2$, $\sigma_a = 0.05$rd. 115

5.6 The Torrance formula: relative intensity versus grazing angle ($\pi/2 - \theta_i$) for $m = 0.9\mu m$, $\Delta\omega = 0.0039$sr, and $\sigma/\lambda = 0.05, 0.1, 0.2, 0.3$. 119

5.7 The Bennett-Porteus formula: relative intensity (for $\Delta\theta = 0.03$sr) versus (a) RMS slope (μm) for $\sigma/\lambda = 0.05, 0.1, 0.2, 0.3$, (b) versus wavelength (μm) for $\sigma = 0.1\mu m$ and $m = 0.01, 0.015, 0.02, 0.03\mu m$. 120

5.8 The Torrance formula: relative intensity (for $\theta = \pi/4$rd and $\Delta\omega = 0.0039$sr) versus (a) RMS slope (μm) for $\sigma/\lambda = 0.05, 0.1, 0.2, 0.3$, (b) versus wavelength (μm) for $\sigma = 0.1\mu m$ and $m = 0.01, 0.015, 0.02, 0.03\mu m$. 120

5.9 The B-K model for the specular direction: (a) relative intensity (i_{spec}) versus grazing angle for $\sigma/\lambda = 0.05, 0.1, 0.2, 0.3$; (b) σ versus λ for $i_{spec} = 0.05, 0.1, 0.15, 0.2$. 123

5.10 Normalized intensity using the modified B-K model (Equation 5.5) versus incidence angle for $m = 1.0, 1.2, .., 1.8$ compared to the Lambertian model (dashed) for two illumination conditions: (a) $\theta_i = \theta_r$ and $\phi_r = \pi$, (b) $\theta_i = \pi/2 - \theta_r$ and $\phi_r = 0$. 127

5.11 Normalized diffuse intensity using the modified B-K model (Equation 5.25) versus incidence angle compared to Lambertian model (dashed) for $\theta_i = \theta_r$ and $\phi_r = \pi$: (a) $T = 0.12\mu m$ and $\sigma = 0.05, ..., 0.2\mu m$, (b) $\sigma = 0.1\mu m$ and $T = 0.1, ..., 0.35\mu m$. 129

5.12 Experimental setup configurations: (a) measuring off-normal specular intensity to estimate σ for slightly-rough dielectrics; (b) measuring two diffuse intensity values to estimate T for slightly-rough or σ/T for very-rough dielectrics; (c) measuring normal specular reflectance to estimate σ for metals; (d) illumination conditions for photometric correction of curved surfaces. 131

5.13 Images of (a-d) metallic surface samples of aluminium, iron, steel and stainless steel scissors under normal incidence ($\theta_i = 0$); (e-g) dielectric surface samples of porcelain plate, plastic plate and glossy paper under off normal incidence ($\theta_i = \theta_r = 14°$). 132

5.14 Samples of (a-b) stone tiles, (c-d) textured wall papers, (e-f) sandpapers, and (g) wooden cube for (1) $\theta_i = 30°$ and (2) $\theta_i = 45°$; (3) enhanced difference images. 134

5.15 Normalized intensity (solid) on sandpaper cylinders (a) type 1 and (b) type 2 versus incidence angle compared to the variants of the B-K model and Lambert's law. 136

5.16 Photometric correction on images of (1) a terra-cotta bear and (2) a sandpaper cylinder: (a) raw images; (b) corrected Lambertian images; (c) difference images. 138

5.17 Photometric correction for (1) the terra-cotta bear and (2) the sandpaper cylinder: (a) Lambertian intensity I_L versus measured intensity I; (b) $(I_L - I)$ versus I. 139

5.18 Removing (b) the Specular components from (a) raw images to obtain (c) diffuse images using the method of Chapter 3; (d) using the modified B-K model to recover Lambertian images; (e) differences between the Lambertian and diffuse images. 140

5.19 Photometric correction for (1) the porcelain bear and (2) the porcelain hand: (a) Lambertian intensity I_L versus diffuse intensity I; (b) $(I_L - I)$ versus I. 141
5.20 Surface normals obtained by applying Lambertian SFS (Chapter 6) to the (a) diffuse and (b) Lambertian images of Figure 5.16; (c) field of differences in needle maps. 142
5.21 Results (as described in Figure 5.20) for the images of Figure 5.18. 142
5.22 Plots for the incidence (a) zenith and (b) azimuth angles (degrees) extracted from the recovered surface normals for (1) the terra-cotta bear and (2) the sandpaper cylinder: the horizontal axis shows the angles corresponding to the diffuse image, while the vertical axis shows those corresponding to the Lambertian image; (c) the difference between the diffuse and Lambertian azimuth angles versus the Lambertian zenith angle. . . . 143
5.23 Re-illumination of the recovered surface normals (Figure 5.20) to generate new images using the modified B-K model for very-rough surfaces for the light-source directions (a-e): $(18^\circ, -90^\circ)$, $(9^\circ, -90^\circ)$, $(0, 0)$, $(9^\circ, 90^\circ)$ and $(18^\circ, 90^\circ)$. 145
5.24 Re-illumination of the recovered surface normals (Figure 5.21) to generate diffuse images using the modified B-K model for slightly-rough surfaces for the light-source directions (a-e): $(40^\circ, -90^\circ)$, $(20^\circ, -90^\circ)$, $(0, 0)$, $(20^\circ, 90^\circ)$ and $(40^\circ, 90^\circ)$. 145
5.25 The ground-truth data of a synthetic vase which is a cylinder with a sinusoid section: (a) the surface height; (b) the surface normal directions. . . . 146
5.26 Surface rendering using (1) the Lambertian model and (2-6) the modified B-K model for very-rough surfaces with an exponential correlation function for the light-source zenith angles (a-d) 0°, 20°, 50°, and 90°; surface slope increases from the second row to the bottom row ($m = 1.0, 1.2, 1.4, 1.6, 1.8$). 147

5.27 Surface rendering using (1) the Lambertian model and (2-6) the modified B-K model for slightly-rough surfaces for the light-source zenith angles (a-d) $0°$, $30°$, $60°$, and $90°$; the RMS roughness is fixed ($\sigma = 0.1\mu m$) while the correlation length increases from the second row to the bottom row ($T = 0.1, 0.15, 0.20, 0.25, 0.30\mu m$). 149

6.1 Smoothing the local Darboux frame: (a) rotating the current surface normal around the irradiance cone by the angle α^k; (b) computing the rotation angle α^k using the unsmoothed and smoothed principal curvature directions, and their projections. 158

6.2 Computing the updated surface normal direction: (a) computing the direction of local surface normal to the smoothed Darboux frame; (b) rotating the closest surface normal onto the irradiance cone to impose the irradiance constraint. 159

6.3 The paraboloid-ellipsoids: ground-truth data (rows 1-2); Worthington-Hancock method (row 3); Darboux smoothing method (row 4). 163

6.4 The crossed paraboloids with two levels of maxima: ground-truth data (rows 1-2); Worthington-Hancock method (row 3); Darboux smoothing method (row 4). 164

6.5 The crossed paraboloids with identical levels of maxima: ground-truth data (rows 1-2); Worthington-Hancock method (row 3); Darboux smoothing method (row 4). 165

6.6 Results for the images of Figures 6.3-6.5: difference vectors between the ground-truth surface normals and those obtained using the Worthington-Hancock method and the Darboux smoothing method (rows 1-2); dot-product images (rows 3-4). 167

6.7 TMS surface normal error versus iteration number obtained using the methods of Worthington-Hancock (dashed curves) and Darboux smoothing (solid curves); panels from left to right correspond to the surfaces shown in Figures 6.3, 6.4 and 6.5 respectively. 168
6.8 A toy duck: surface normals obtained using the two methods (row 1); difference between the two sets of surface normals; minimum curvature directions (row 2). 169
6.9 A terra-cotta watering-can: surface normals obtained using the two methods (row 1); difference between surface normals; minimum curvature directions (row 2). 170
6.10 A terra-cotta teapot: surface normals obtained using the two methods (row 1); difference between surface normals; minimum curvature directions (row 2). 171
6.11 Results for the images of Figures 6.8-6.10: (a-b) the shape-index, (c-d) the curvedness and (e) dot-product images for the surface normals obtained using the Worthington-Hancock method and the Darboux smoothing method. 172
6.12 A synthetic surface containing 16 ellipsoids on top of a large paraboloid. . 173
6.13 Row 1: fields of surface normals using the uncompensated (left) and compensated (right) Hessian in the D-S method; row 2: differences with ground-truth. 173
6.14 Row 1: recovered azimuth angles using the D-S method with the uncompensated (left) and compensated (right) Hessian versus ground-truth azimuth angles; row 2: difference between the recovered and ground-truth azimuth angles versus zenith angles. 174
6.15 The shape-index (row 1) and the curvedness (row 2) images of the ground-truth (left) and those using the uncompensated (middle) and compensated (right) Hessian. 175

6.16 TMS surface normal error versus iteration number obtained using the uncompensated (dashed curve) and compensated (solid curve) Hessian methods. 175
6.17 A terra-cotta coffee-pot: fields of surface normals (row 2) and minimum curvature directions (row 3) obtained using the uncompensated (left) and compensated (middle) Hessian in the D-S method; (right) differences between the two sets of vectors. 176

Acronyms

SFS	Shape-From-Shading
RMS	Root-Mean-Square
TMS	Total-Mean-Square
D-S	Darboux Smoothing
B-K	Beckmann-Kirchhoff
W-H	Worthington-Hancock
T-S	Torrance-Sparrow
O-N	Oren-Nayar
G-T	Ground-Truth

Symbols

$\vec{L}$	light-source direction
$\vec{V}$	viewing direction
$\vec{N}$	surface normal direction
Ψ	the angle between $\vec{L}$ and $\vec{V}$
E	grey-level image
L_r	reflected radiance
I	optical intensity
F	Beckmann's geometrical function
f	Fresnel coefficient
n	index of refraction
λ	wavelength
$\vec{N}_S$	specular spike direction
$\vec{N}_L$	Lambertian surface normal
$\vec{M}$	posterior-mean surface normal (Chapter 3)
$\vec{M}$	principal curvature direction (Chapter 6)
σ	surface RMS roughness
T	surface correlation length
m	surface RMS slope
H	Hessian matrix
ϕ	the shape-index
w	weigth

Acknowledgements

I am pleased to express my gratitude to Professor Edwin Hancock for his excellent supervision of my PhD research. In particular, I am impressed by his great attitude and his ongoing scientific activities within the computer vision community. I am also grateful to my assessor Dr Richard Wilson for his constructive advice on my scientific reports.

When it comes to my family, the first person I would like to thank is my wife whose kind support during the research program must be praised, and her patience during the past few years must be appreciated. Next, I feel very proud and grateful to my parents who have been wisely encouraging me to continue learning and teaching during my primary, secondary and higher educations. Although they ask for nothing in return, at this stage, it is my duty to show them my special respects for all their invaluable efforts.

My colleagues at the University of York deserve to be acknowledged for exchanging their knowledge and experience with me and for providing sample computer source codes and technical advice. In particular, I would like to thank those colleagues who presented a number of my papers on my behalf at various conferences which I could not attend.

Finally, my sponsor, the Ministry of Science, Research and Technology (MSRT) of the Iranian government, who granted me both the scholarship and reasonable financial support should be named here with respects and regards.

Author’s Declaration

I declare that all the work presented in this thesis is solely my own, except where attributed and cited to another author. Some of the material presented in different chapters of this thesis have previously appeared in the publications which are listed separately.

List of Publications

Journal Papers

- Ragheb, H. and Hancock, E. R., Darboux Smoothing for Shape-from-Shading, *Pattern Recognition Letters*, Vol. 24, Nos. 1-3, 2003, pp. 579-595.
- Ragheb, H. and Hancock, E. R., A Probabilistic Framework for Specular Shape-from-Shading, *Pattern Recognition*, Vol. 36, No. 2, 2003, pp. 407-427.
- Ragheb, H. and Hancock, E. R., Physical Reflectance Models for Rough Surface Analysis, *IEEE Trans. on Pattern Analysis and Machine Intelligence*, submitted paper.
- Ragheb, H. and Hancock, E. R., Radiance Correction for Shape-from-Shading, *IEEE Trans. on Pattern Analysis and Machine Intelligence*, submitted paper.

Conference and Workshop Papers

- Ragheb, H. and Hancock, E. R., Shape-from-Shading using Darboux Smoothing, *The Ninth International Conference on Computer Analysis of Images and Patterns*, September 5-7, 2001, Warsaw, Poland, LNCS 2124, pp. 657-667, *The Winner of the Young Scientist Award*.
- Ragheb, H. and Hancock, E. R., Separating Lambertian and Specular Reflectance Components using Iterated Conditional Modes, *The Twelfth British Machine Vision Conference*, September 10-13, 2001, Manchester, UK, pp. 541-552.
- Ragheb, H. and Hancock, E. R., Improved Shape-from-shading Using Darboux Smoothing, *The Eighth International Conference on Image Processing*, October 7-10, 2001, Thessaloniki, Greece, pp. 817-820.
- Ragheb, H. and Hancock, E. R., Specularity Subtraction for Shape-from-Shading, *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, December 9-14, 2001, Kauai Marriott, Hawaii, Technical Sketches CD.

- Ragheb, H. and Hancock, E. R., Lambertian Correction for Rough and Shiny Surfaces, *The Ninth International Conference on Image Processing*, September 22-25, 2002, Rochester, New York, USA, Part 2, pp. 553-556.

- Ragheb, H. and Hancock, E. R., Shape-from-shading using viewpoint-invariant principal curvatures, *The Ninth International Conference on Image Processing*, September 22-25, 2002, Rochester, New York, USA, Part 2, pp. 577-580.

- Ragheb, H. and Hancock, E. R., Highlight Removal using Shape-from-Shading, *The Seventh European Conference on Computer Vision*, May 27-31, 2002, Copenhagen, Denmark, LNCS 2351, Part 2, pp. 626-641, Copenhagen, Denmark.

- Ragheb, H. and Hancock, E. R., A Probabilistic Framework for Specular Shape-from-Shading, *The Sixteenth International Conference on Pattern Recognition*, August 11-15, 2002, Qubec, Canada, pp. 513-516.

- Ragheb, H. and Hancock, E. R., Shape-from-Shading for Highlighted Surfaces, *The Fourth International Workshop on Statistical Techniques in Pattern Recognition*, August 6-9, 2002, Windsor, Canada, LNCS 2396, pp. 576-586.

- Ragheb, H. and Hancock, E. R., Diffuse Reflectance Models for Rough Surfaces: A Geometrical Study for Shape-from-Shading, *The Second International Workshop on Texture Analysis and Synthesis*, June 1, 2002, Copenhagen, Denmark, pp. 113-116.

- Ragheb, H. and Hancock, E. R., Non-Lambertian Shape-from-Shading using Iterated Conditional Modes, *The Sixth International Statistics Conference*, August 26-28, 2002, Tehran, Iran, pp. 198-219.

- Ragheb, H. and Hancock, E. R., Rough Surface Analysis using Kirchhoff Theory, *The Fourteenth British Machine Vision Conference*, September 9-11, 2003, Norwich, UK, Part 1, pp. 191-200.

- Ragheb, H. and Hancock, E. R., Improving Shape Recovery by Estimating Properties of Slightly-Rough Surfaces, *The Twelfth International Conference on Image Analysis and Processing,* September 17-19, 2003, Mantova, Italy, pp. 32-37.

- Ragheb, H. and Hancock, E. R., Rough Surface Correction and Re-illumination using the Modified Beckmann Model, *The Tenth International Conference on Computer Analysis of Images and Patterns*, August 25-27, 2003, Groningen, The Netherlands, LNCS 2756, pp. 98-106 .

- Ragheb, H. and Hancock, E. R., Estimating Surface Characteristics using Physical Reflectance Models, *IEEE Computer Society Conference on Computer Vision and Pattern Recognition,* June 16-22, 2003, Madison, USA, Part 2, pp. 177-184.

- Ragheb, H. and Hancock, E. R., Rough Surface Estimation using the Kirchhoff Model, *The Thirteenth Scandinavian Conference on Image Analysis,* June 29 - July 2, 2003, Gteborg, Sweden, LNCS 2749, pp. 477-484.

- Ragheb, H., Robles-Kelly, A. and Hancock, E. R., Surface Radiance Data Against Reflectance Models, *The Ninth Computer Vision Winter Workshop,* February 4-6, 2004, Piran, Slovenia, pp. 119-128.

- Ragheb, H. and Hancock, E. R., Surface Radiance: Empirical Data Against Model Predictions, *The Eleventh International Conference on Image Processing*, October 24-27, 2004, Suntec City, Singapore, to appear.

Chapter 1

Introduction

The recovery of surface shape from shading patterns has been a topic of sustained activity in the computer vision literature for almost three decades (Horn, 1970). There is strong psychophysical evidence for its role in surface perception and recognition (Koenderink et al., 1992). The problem of recovering surface normal and height data from variations in surface brightness is an under-constrained one which relies on the use of simplifying assumptions and constraints to be rendered tractable. One of the most limiting of these is that the reflectance of the surface follows Lambert's law. This originates from the fact that the reflectance properties of both rough surfaces (where there is a significant variation in local surface relief) and shiny surfaces (where there is significant microfacet reflection) depart significantly from the predictions of Lambert's law (which applies to surfaces which are smooth and matte). Specifically, the main departures occur due to specularities, subsurface scattering and rough surface scattering.

The main aim in this thesis is to develop methods for handling non-Lambertian reflectance in shape-from-shading (SFS). To do this we draw on models of specular and diffuse reflectance from rough and smooth surfaces. The main contributions are:

- To develop a probabilistic method for specularity subtraction. Here we use the iterative conditional modes algorithm of Besag to separate the Lambertian and specular reflectance components. The specular component is modelled using the Torrance and Sparrow model.

- To explore how the phenomenological reflectance models of Oren-Nayar, Wolff and Wolff-Nayar-Oren can be used to correct large-angle reflectance effects for rough and smooth surfaces. We show how the reflectance models can be used to recover corrected Lambertian reflectance.

- To explore how the wave scattering theory developed by Beckmann can be used for rough surface analysis. Here we explore how different variants of the Beckmann model can be applied to surfaces of varying degrees of roughness. We illustrate how the different variants of the Beckmann model can be used to estimate surface roughness parameters and recover Lambertian reflectance.

By subtracting specularities and recovering corrected Lambertian reflectance, we demonstrate that more accurate surface analysis using SFS is possible. As a subsidiary contribution, we present a new method for surface normal smoothing in SFS. Here we demonstrate how surface normal directions can be smoothed by aligning the Darboux frames of adjacent image locations. The main advantage of this method is to improve the recovery of ridge and ravine structures on curved surfaces.

In this chapter we describe the organization of the different chapters of the thesis and summaries the main assumptions made throughout the thesis. The research presented in the thesis is divided into three major parts. In the first part, we tackle the problem of removing specular highlights from shiny surfaces. In the second part, we focus on how a variety of both phenomenological and physical reflectance models can be utilized to perform photometric correction on real-world images of surfaces with different scales of roughness. Finally, in the third part, the use of curvature information for developing a new smoothness constraint for SFS, called Darboux smoothing, is presented. It should be noted that, each individual piece of research presented in the thesis represents a contribution towards the final goal of improved SFS. It is worth mentioning that the order of the chapters is not the same as that in which we performed the various pieces of research. Instead, we first performed the research on separating specular and non-specular components. We then assumed that the non-specular component follows Lambert's law,

and developed a SFS algorithm using Darboux smoothing. Finally, we investigated the use of different reflectance models for rough surfaces where Lambert's law fails to distinguish between surfaces with different scales of roughness. However, it is believed that the presented order of chapters results in a more coherent thesis structure.

In Chapter 2, we review the related literature from computer vision and optical physics. Specifically, we focus on existing physics-based, phenomenological and empirical reflectance models for the specular and diffuse components. We also consider whether these models have been reformulated, modified or used to improve the quality of surface normal directions delivered by SFS.

In Chapter 3, we develop a SFS algorithm which can deal with images of specular surfaces based on a geometric-probabilistic framework which allows both Lambertian and specular reflectance models to be utilized. We also show how by reconstructing the specular intensity component we can remove specular highlights from the raw image. The separated diffuse image can be used as input to Lambertian SFS methods. We make use of the Torrance and Sparrow model (Torrance and Sparrow, 1967) to reconstruct the specular intensity from the estimated surface normal directions. We also use the Cook and Torrance model (Cook and Torrance, 1981), and, the Healey approximation (Healey and Binford, 1987) where we investigate different distribution functions for the microfacet normal directions. The basic method of Worthington and Hancock for smoothing surface normal directions (Worthington and Hancock, 1999) is used to perform SFS on the Lambertian component. Finally, the "iterated conditional modes" of Besag (Besag, 1986) provides a simple and effective probabilistic framework for computing the proportions of the two reflectance components, i.e. the specular and Lambertian components.

Our specular SFS method is expected to result in enhanced shape recovery and advantages over recent SFS methods. The novelty of the work lays in the fact that although color photometric stereo have been used for specularity removal, there are few examples of doing this with single grey-level images. One reason for the lack of any similar approach to ours in the literature may lay in the fact that it relies on an extension of the geometrical framework proposed recently by Worthington and Hancock (Worthington and

Hancock, 1999) for iterative SFS. By looking at the problem from this point of view, we have developed our geometric-based SFS method which can deal with specular images.

The second topic addressed in the thesis is that of non-specular reflectance from rough and shiny surfaces. Although the phenomenological reflectance models (Oren and Nayar, 1995; Wolff, 1994a; Wolff et al., 1998) provide a more accurate account of reflectance from rough and shiny surfaces, they do not prove tractable for estimating surface normal directions in SFS or other problems involving the analysis of surface characteristics or shape. There are two reasons for this. Firstly, the models depend on both the direction of the incident light and the viewer. Secondly, even when these two directions are identical, some models do not become analytically tractable because of their complex formulations. Nonetheless, as we shall show in Chapter 4, they do provide a route by which the Lambertian component can be recovered and used for the purposes of shape recovery. The method proposed is a photometric correction process. This contribution is also novel because there has been little effort aimed at using these reflectance models for curved surface analysis.

Since phenomenological reflectance models are not based purely on physics, they do not describe well physical surface properties, and so, can not be used to estimate meaningful rough surface physical parameters. Hence, in Chapter 5, we investigate the applicability of some physics-based models for the reflectance modelling for rough surfaces. The models studied here are based on wave scattering theory and commence by the Kirchhoff integral. Specifically, we use wave scattering theory to develop new surface parameter estimation techniques using different variants of the Beckmann-Kirchhoff model (Beckmann and Spizzichino, 1963). We also make use of the modification of Vernold and Harvey (Vernold and Harvey, 1998) to adapt the Beckmann model for computer vision applications. There has been previous research on adapting the Kirchhoff theory for computer graphics applications. However, in our knowledge, we are the first to adapt the Beckmann scattering model based on the Kirchhoff theory for surface shape analysis from digitized images. Moreover, similar techniques for wave scattering from conducting surfaces are presented using models derived from Davies theory (Bennett and Porteus,

1961; Torrance, 1964). We then show that the photometric correction process described in Chapter 4 is also applicable to different variants of the modified Beckmann model for diffuse intensity. This enables us to use the wave scattering models to recover corrected Lambertian reflectance, which can then be used as input to SFS.

Even for images which follow Lambert's law, SFS is an under-constrained problem which requires constraints on consistent surface normal directions to be rendered tractable. In Chapter 6, we introduce a new method which we refer to as Darboux smoothing which is used in conjunction with our Lambertian SFS algorithm to recover improved surface normal directions. The method exploits curvature information in the smoothing process using differential geometry (de Carmo, 1976). Rather than just smoothing the field of surface normals, we smooth the orientations of the local Darboux frames instead. Specifically, the shape-index quantity of Koenderink and van Doorn (Koenderink and van Doorn, 1992) is employed to compute the local weights for controlling the smoothing of local principal curvature directions. The corrected Lambertian reflectance component extracted using the models described in Chapters 3, 4 and 5 can be used as input to the SFS method described in Chapter 6.

Finally, in the concluding chapter, we review the results of the whole thesis and provide some recommendations for the future research. One obvious line of future investigation is to use surface normal directions delivered by our SFS method together with other available clues for the purpose of object and face recognition. However, this is beyond the scope of the thesis.

Chapter 2

Literature Review

In this chapter, we review the literature related to the topic of the thesis. The review commences by considering the general problem of shape-from-shading (SFS), and then focuses in some detail on methods for modelling non-Lambertian reflectance. The structure of the chapter is as follows. In Section 2.1, we discuss classical and recent techniques for SFS, and discuss their shortcomings. One of the most important of these is that they fail to deal with non-Lambertian reflectance. Hence, in Section 2.2 we focus in some detail on non-Lambertian reflectance models and methods for the analysis of reflectance from different classes of rough surface. There are a number of contributions which are not directly related to the work presented in this thesis, but present important research related to the general problem of SFS. Hence, in Section 2.3., we review some of these publications briefly. Specifically, we examine some of the related research on photometric stereo including techniques for separating specular and diffuse reflectance components. Then, we discuss the literature which stresses the importance of illumination in SFS. Finally, we look at some recent methods for analyzing textured surfaces (including the BRDF approach). This work has acquired recent topicality since it is important in bridging the gap between shape-from-shading and shape-from-texture. Finally, in Section 2.4, we use the review of the literature to motivate the contributions presented in the remaining chapters of the thesis.

2.1 Shape from Shading

Shape-from-shading is the problem of recovering surface orientation and hence surface height using the variations in shading across a curved object. This problem is usually posed as that of recovering the field of surface normals that satisfy the image irradiance equation. With the field of surface normals to hand, then the height function can be reconstructed by performing surface integration (Woodham, 1994).

Some of the pioneering work in the area is reported in the PhD theses by Horn (Horn, 1970) and by Krakauer (Krakauer, 1971). Early developments in the field are well documented by Horn in the chapter "Height and Gradient from Shading" which appears in the collection of papers entitled Shape-from-Shading (Horn and Brooks, 1989). There have been a number of subsequent attempts to improve and develop this work. Okatani and Deguchi (Okatani and Deguchi, 2000) have divided SFS algorithms into three groups. These algorithms make use of image information in a global (Horn and Brooks, 1989; Horn, 1990; Oliensis, 1991b; Oliensis, 1991a; Kimmel and Bruckstein, 1995a; Wei and Hirzinger, 1997), semi-local (Bruckstein, 1988; Bichsel and Pentland, 1992; Oliensis and Dupuis, 1993; Brooks and Chojnacki, 1994; Kimmel and Bruckstein, 1995b; Kimmel et al., 1995; Prados et al., 2002) or local (Pentland, 1984; Lee and Rosenfeld, 1985; Sara, 1994) manner.

Zhang et al. (Zhang et al., 1999), on the other hand, have classified SFS techniques into four categories. These are minimization approaches (Ikeuchi and Horn, 1981; Brooks and Horn, 1985; Frankot and Chellappa, 1988; Horn, 1990; Malik, 1989; Szeliski, 1991; Zheng and Chellappa, 1991; Lee and Kuo, 1993; Leclerc and Bobick, 1991; Vega and Yang, 1993), propagation approaches (Horn, 1970; Prados et al., 2002; Dupuis and Oliensis, 1992; Kimmel and Bruckstein, 1992; Bichsel and Pentland, 1992), local approaches (Pentland, 1984; Lee and Rosenfeld, 1985), and linear approaches (Pentland, 1988; Tsai and Shah, 1994). However, there are SFS methods which fall outside these four classes. However, our aim is not to present a new classification of SFS schemes in this thesis.

Unfortunately, despite considerable effort over the past three decades, reliable shape

recovery from shading information has proved an elusive goal. The reasons for this are twofold. Firstly, the recovery of surface orientation from the image irradiance equation is an under-constrained process which requires the provision of boundary conditions and constraints on surface smoothness to be rendered tractable. Secondly, real-world imagery rarely satisfies these constraints.

In fact, measured against current progress, some of the earliest work on the subject of SFS was ambitious in its goals, allowing for both general reflectance models and perspective geometry. Although the bulk of the recent work in SFS restricts itself to a Lambertian model, the need for non-Lambertian reflectance models has been appreciated from the outset. For instance, in work from photoclinometry Rindfleish (Rindfleish, 1966) has used the non-Lambertian reflectance models of Hapke, Minnaert and Fesenkov (Hapke, 1963; Minnaert, 1961; Fesenkov, 1962) to recover the slopes of lunar maria. Moreover, the reflectance map of Horn (Horn, 1981) is not restricted to Lambertian reflectance. One area where non-Lambertian reflectance models are of critical importance, and have been used with some success, is in radar SFS. This is a process of recovering surface gradient and height from radar returns, and has proved to be an important tool in the analysis of Megallan radar images of Venus. Here, Frankot and Chellappa (Frankot and Chellappa, 1988) have used an empirical radar reflectance model in conjunction with a Fourier domain method for reconstructing height under integrability constraints. In a recent paper, Bors et al. (Bors et al., 2003) have developed a statistical method for radar SFS, using a digital elevation map to estimate the radar reflectance function.

2.1.1 Curvature Consistency

Curvature consistency constraints have been demonstrated by several authors to improve the quality of surface topography information recoverable using SFS. There are several contributions in the literature which deserve special mention. Frankot and Chellappa (Frankot and Chellappa, 1987) have focused on the differential geometry of the recovered surface and have developed a Fourier domain approach for imposing integrability con-

straints. Ferrie and Lagrarde (Ferrie and Lagarde, 1990) have used the Darboux-frame smoothing method of Sander and Zucker (Sander and Zucker, 1990) as a post-processing step in SFS to impose constraints from differential geometry on the recovered needle map. These constraints are provided by the principal curvature directions. However, there are limitations to the use of least-squares smoothing of the Darboux frames. These relate to the fact that the principal curvature direction is undefined at umbilic surface points and that the direction is discontinuous at saddles. Recently, Worthington and Hancock (Worthington and Hancock, 1999) have used error kernels suggested by robust statistics for the purposes of smoothing. Curvature consistency constraints on the recovered surface normals are imposed in their method using the shape-index (an angular measure of local surface topography) of Koenderink and van Doorn (Koenderink and van Doorn, 1992).

However, one of the problems that limits these curvature sensitive smoothing methods are the effects of foreshortening for highly inclined surfaces. The main problem is that the Hessian matrix is not invariant under changes in surface inclination. To overcome this problem, Woodham (Woodham, 1994) introduced the viewpoint-invariant Hessian matrix as an alternative to the viewer-centered Hessian matrix.

2.2 Reflectance Modelling

Reflectance modelling is a task of pivotal importance in the analysis of image data by computer. For instance, in computer graphics it is necessary for generating realistic images of synthetic objects. In computer vision, on the other hand, reflectance models form the basis of shape analysis techniques such as SFS and photometric stereo, and may also be used to estimate the physical properties of materials from passively sensed image data (van Ginneken et al., 1998).

Lambert's law is the simplest model for diffuse reflection from a matte surface. It has been widely used for different applications in computer vision, and specifically, to recover shape from shading patterns. According to Lambert's law, for a matte surface of constant albedo illuminated with a single collimated light-source, the observed intensity is

independent of the viewing direction. The observed intensity depends only on the surface albedo and the quantity of absorbed light, and this in turn is proportional to the cosine of the incidence angle between the light-source direction and the local surface normal.

Although the bulk of existing work on SFS assumes that the reflectance of the surface under study follows Lambert's law, the quest for methods that can deal with non-Lambertian reflectance is a long standing one (Belhumeur and Kriegman, 1996). There are a number of situations in which significant departures from Lambert's law both for rough and shiny surfaces have been reported. The most significant departures occur at the locations of specularities and also close to the object occluding boundary (Wolff et al., 1998). It should be stressed that the basis of Lambert's law is purely phenomenological, and the reflectance function does not attempt to model the physical processes of light scattering. In this section we review some non-Lambertian reflectance models.

2.2.1 What Roughness Means

Before discussing the different reflectance models, it is important to understand the physical meaning of surface roughness and its use as a surface characteristic. Surface roughness is the critical parameter in determining departures from Lambert's law. Surface roughness can be characterized in a number of ways. For very-rough surfaces, one approach is to use a model which describes the distribution of surface wall cavity angles (Oren and Nayar, 1995). For slightly-rough surfaces, i.e. ones that present a shiny appearance, roughness can be modelled using the angular distribution of microfacets (Torrance and Sparrow, 1967). An alternative that can capture both effects is to describe the roughness phenomenon using the variance and the correlation length of variations in the surface height distribution (van Ginneken et al., 1998).

Bennett and Mattsson (Bennett and Mattsson, 1999) have defined the roughness as a measure of the statistical variation in the distribution of topographic relief of a surface. They observe that the most frequently used statistical measure of roughness is the root-mean-square (RMS) height σ. Although there is conflicting evidence, the height distribu-

tion is frequently assumed to be Gaussian (Ogilvy, 1991). Unfortunately, the specification of a height distribution and the RMS roughness is insufficient to discriminate between surfaces with different length scales. Such surfaces may, however, be distinguished on the basis of their correlation length T (Ogilvy, 1991). It is frequently assumed that the correlation function is also Gaussian. However in practice, the exponential correlation function is sometimes found to give a better fit to the measured surface data (Ogilvy, 1991). Length and height effects may be jointly characterized using the RMS slope m. The statistical distribution of the slope parameter for real surfaces can be estimated using the distribution of the difference between the measured heights of adjacent points divided by the data sampling interval. The RMS slope can then be deduced from this distribution.

RMS roughness can be estimated directly from surface-profile measurements, or it can be calculated from a scattering measurement using a theoretical reflectance model. Although they are easier to make, scattering measurements yield only a statistical average of the surface roughness (Bennett and Mattsson, 1999). The determination of correlation length depends on the form of the correlation function and is relatively straightforward when the correlation function is either Gaussian or exponential. However, for more complicated correlation functions the determination of correlation length is less tractable.

The effect of variable surface roughness is to alter the relative contributions of specular and diffuse reflectance to the total reflectance (He et al., 1991). For instance, a perfectly smooth surface has purely specular reflectance and hence acts as a mirror by reflecting the light in the incidence plane so that the incidence and reflectance angles are equal. Rough specular surfaces, on the other hand, reflect light into a lobe around the specular direction. In fact, the specular reflection is a very sharp function of incidence and reflectance angles and is nearly zero for all scattering directions except a very narrow range around the specular direction (or specular spike direction). Nayar et al. (Nayar et al., 1991) referred to this region as the specular lobe. As first observed by Lambert, under diffuse reflection, matte surfaces without macroscopic roughness appear equally bright from all viewing directions. For a rough matte surface, however, the surface appears brighter as the viewing direction approaches the illumination direction. The effect of roughness is also to change

the total reflectance, but this effect is weak in most cases and only becomes significant for particular viewing and illumination configurations (Beckmann, 1967).

2.2.2 Specular Reflectance Models

The modelling of specular reflectance and specular highlights has attracted considerable attention in the computer vision and computer graphics communities. There are a number of ways in which specular reflectance can be modelled. For instance, in the graphics community some purely empirical models have been developed with the goal of developing computationally efficient tools for the purposes of performing realistic surface synthesis. Perhaps the most widely used of these is the Phong model (Phong, 1975) which is a parameterized continuous function that can be used to represent specular reflectance. The model was first used for rendering images of objects in computer graphics. The simplest model for specular reflection is the double-delta model used by Ikeuchi (Ikeuchi, 1981) to determine the shape of specular surfaces by photometric stereo. However, physically motivated models have proved to be more realistic and reliable, particularly for applications outside the graphics domain.

The Torrance-Sparrow model (Torrance and Sparrow, 1967), which is perhaps the most commonly used model, represents one of the first attempts towards the modelling of specular intensity. The model captures the underlying physics of scattering by assuming a microfacet structure for the rough surface. The surface is modelled as a collection of specular V cavities whose surface normals are assumed to be normally distributed. Each cavity consists of two planar facets and the Fresnel theory of the air-surface boundary individually applies to each facet. The roughness of the surface is specified using a distribution function for the facet slopes. However, this model is a geometrical one and although it possesses a simpler mathematical form than physical models, it is only applicable when the wavelength of the incident light is small compared to the dimensions of the surface irregularities. Cook and Torrance (Cook and Torrance, 1981) modified the Torrance and Sparrow model and used it to render images. Commencing from the Torrance-Sparrow

model, Healey and Binford (Healey and Binford, 1987) derived a simplified model by using the Gaussian distribution as the facet orientation function and considering the other components as constants. By adopting this simplified model, they implemented a shape from specularity system. Their main contribution is to derive important relationships between the properties of a specular feature in an image and the local properties of the corresponding surface. Recently, Angelopoulou and Poger (Angelopoulou and Poger, 2003) have looked critically at the Cook and Torrance model for specular intensity (Cook and Torrance, 1981) and found that the modelling error increases dramatically as the incidence angle increases.

Several authors have looked critically at the physics underlying specular reflectance. For instance, in the reflectance model described by He et al. (He et al., 1991), the specular component depends on the effective roughness, while the directional dependance of the diffuse component is governed by both the correlation function of the surface and the effective roughness. Also, Barrick (Barrick, 1968) derived an expression for the average number of specular points on a Gaussian surface. This surface model has been used by van Ginneken et al. (van Ginneken et al., 1998) to obtain a reflection model for both the specular and diffuse components.

Several authors have used specularities for the purpose of shape analysis. For instance, in a comprehensive treatment of specular SFS, Brelstaff and Blake (Blake and Brelstaff, 1988; Brelstaff and Blake, 1988) have analyzed the geometric constraints provided by specularities and have shown how to detect specularities using Lambertian irradiance constraints. Clark and Yuille (Clark and Yuille, 1990) considered the contributions of both Lambertian and specular reflectance components to develop an algebraic approach for recovering shape information from shiny images. Drawing on psychophysics, Blake and Bulthoff (Blake and Bulthoff, 1991) have developed a computational model of the shape information available to a moving observer from the apparent movement of specularities. More recently, Drbohlav and Sara (Drbohlav and Sara, 2002) have presented a method which uses specularities to reduce the ambiguity in photometric stereo with uncalibrated light directions and intensities.

Specular and diffuse reflectance separation is an important subject in computer vision. For instance, in a series of recent papers, Lin and Lee (Lin and Lee, 1999b; Lin and Lee, 1999a; Lin and Lee, 2000) have shown how specular reflection due to multiple light-sources can be located and separated in multi-band imagery. Later, Lin and Shum (Lin and Shum, 2001) have developed a method for separating the diffuse and specular reflectance components in color images using two photometric images without calibrated lighting. More recently, Lin et al. (Lin et al., 2002) have proposed a method for estimating both the separation and the true depth of specular reflection. The approach integrates color analysis and multibaseline stereo in a synergistic manner. Also, Swaminathan et al. (Swaminathan et al., 2002) have proposed a technique to extract specularities as a separate layer and have experimented using image sequences of a complex scene.

Drawing on polarization information, Nayar et al. (Nayar et al., 1993) have used polaroid filters to detect specular reflection. In a very recent paper, Miyazaki et al. (Miyazaki et al., 2004) have shown how the surface shape of transparent objects can be estimated using polarization information. They have used the degree of polarization at each surface point to find the surface normal direction. To solve the ambiguity problem in polarization data, they have used two different images from the object.

2.2.3 Physics-Based Reflectance Models

Although an increasing number of papers in the computer vision literature are motivated by the need to incorporate more principled reflectance models, they fall short of developing purely physics-based models that account for directly measurable roughness parameters. In this section we provide a brief review of some physics-based models for the scattering of light from rough surfaces. The mathematical detail of these models is presented in Chapter 5. One of the most straightforward approaches is to develop a model based on geometrical optics. Underlying the approach is the assumption that the scattering surface consists of planar facets having lateral dimensions that are large compared to the wavelength of the incident light and which reflect like plane mirrors (Elson et al.,

1979). Each facet has a surface normal and there will be a statistical distribution of the directions of the surface normals. Thus, according to geometrical optics, scattering from a rough surface consists of summing the contributions from each tilted facet of the surface. In order to analyze a surface using this approach, one needs to have a characterization of the surface to hand. This can be either a model of the shape and angular distribution of the surface facets or an equivalent statistical characterization (Elson et al., 1979). For instance, Beckmann (Beckmann and Spizzichino, 1963) developed a Markov chain model to account for the surface statistics. However, in general such characterizations are not easily obtained. Also, geometrical optics breaks down when the wavelength of the light and the facet dimensions are comparable. It hence fails to describe the effects of interference, diffraction, and polarization. For these reasons, geometrical optics is rarely used for optical scattering problems (Elson et al., 1979).

Geometric optics involves ray-tracing and is diffraction limited, whereas wave scattering theory models the interactions between the wave and the surface. The scalar version of wave scattering accounts only for the distribution of intensity, and does not account for more complex effects such as polarization. Scalar wave scattering models have been applied both to surfaces whose roughness is much less than the wavelength of light (Bennett and Porteus, 1961; Davies, 1954; Depew and Weir, 1971) and also to those whose roughness is much greater than the wavelength of light (Chandley, 1976; Davies, 1954). The starting point for the scalar wave theory treatment of scattering from rough surfaces is the Helmholtz-Kirchhoff diffraction integral (Bruce and Dainty, 1991). To overcome some problems involved in solving this integral, an approximation known as the Kirchhoff boundary condition is made. The Beckmann formulation of the scalar Kirchhoff theory (referred to here as the Beckmann-Kirchhoff or B-K theory) is the most widely used in the study of wave scattering from rough surfaces (Beckmann and Spizzichino, 1963; Beckmann, 1967).

Despite its widespread use, the B-K theory is based on a paraxial (small-angle) assumption that limits its ability to accurately account for both wide-angle scattering and large angles of incidence (Harvey et al., 1999; O'Donnell and Mendez, 1987). Recently,

Vernold and Harvey (Vernold and Harvey, 1998) have modified the B-K theory to overcome this limitation. Their modification involves replacing the geometrical form factor of the B-K theory by a simpler Lambertian form. This extends the utility of the theory to large angles of both incidence and scatter. The modification is based on empirical and phenomenological arguments. However, this does not diminish the pragmatic interest in the modified B-K model since it includes some effects (corresponding to non-paraxial angles) which are not included in the original B-K model. Later, Harvey et al. (Harvey et al., 1999) have critically analyzed this problem and have developed a linear systems approach to wide-angle scalar diffraction theory.

On the other hand, Kirchhoff theory is a single scattering theory and does not account for multiple scattering. When multiple scattering occurs, energy is not transferred directly from the incident wave into the scattered wave, but instead becomes "trapped" for a limited time within the hills and troughs of the rough surface (Ogilvy, 1991). Hence, Kirchhoff theory underestimates the scattered intensity at large surface slopes. Authors from physical optics have attempted to improve the predictions of scattered intensity delivered by Kirchhoff theory. For instance, Bruce and Dainty (Bruce and Dainty, 1991) refined the B-K theory to include multiple scattering and shadowing effects. Recently, Caron et al. (Caron et al., 2002) have looked critically at different problems of Kirchhoff theory. They have attempted to model the multiple scattering effect using the energy conservation rate. This is the ratio of the scattered energy to the incident energy. A second limitation of the Kirchhoff theory identified by them is that for rough surfaces under oblique incidence there are problems due to surface self shadowing. To address this problem, they have introduced a limit for angles of incidence under which the predictions of the Kirchhoff model are reliable.

It is interesting to note that Kirchhoff theory has been used to develop physically motivated reflectance models for computer graphics applications. For instance, He et al. (He et al., 1991) proposed a comprehensive model that incorporates complex factors including surface statistics, sub-layer scattering, and polarization. They exploited an improved representation of the Kirchhoff integral (Stogryn, 1967) for surface synthesis. As a re-

sult, their model incorporates a more detailed model of the effects of surface height and slope. More recently, Stam (Stam, 1999) has applied the Kirchhoff approach to include diffraction effects in surface rendering.

A practical problem with the Beckmann model (Beckmann and Spizzichino, 1963) is that for different scales of surface roughness, one should use different variants of the model. This problem has been overcome by Sun (Sun, 2003) who has developed a "one-bounce" reflectance model using physical parameters which is applicable to both smooth and rough surfaces. The model is derived using the Fresnel coefficient, a self-shadowing factor and a probability function for surface orientation. It is interesting that the model generates a sharp specular highlight for smooth surfaces and gives a diffuse behavior for rough surfaces.

Physical optics has lead to a number of alternative approaches to the problem of modelling specular reflection from rough surfaces. For instance, the Beckmann distribution function can be used to model the angular distribution associated with the reflected light in the proximity of specularities (Beckmann and Spizzichino, 1963). It can describe reflection from surfaces with a wide range of roughness including very smooth surfaces. Nayar et al. (Nayar et al., 1991) showed that the Beckmann model predicts both the specular lobe and the specular spike components while the Torrance-Sparrow model only predicts the specular lobe component. They also showed that in the small wavelength limit, the Torrance-Sparrow model is in fact a more elaborated version of the Beckmann model.

There are alternative reflectance models which apply for rough conducting surfaces only. Some of these models are based on the theory of reflection of electromagnetic radiation from rough surfaces which was derived by Davies (Davies, 1954). The surface is represented by a statistical model. This theory was originally developed to account for the scattering of radar waves from rough water surfaces. Several authors have adapted Davies theory for applications in the optical region. For instance, there are two refinements of the Davies model which apply under different restrictions concerning the reflectance geometry of the specular direction. The first of these is due to Bennett and Porteus (Bennett and Porteus, 1961) and applies for normal incidence and reflectance. The second is due

to Torrance (Torrance, 1964) and applies when the incident radiation is off normal.

Bennett and Porteus (Bennett and Porteus, 1961) also described a technique for determining the surface roughness of opaque materials using an optical method. Meanwhile, they gave attention to the strong inverse dependency of the diffuse component on wavelength. In particular, they noted that, at wavelengths which are long compared to the roughness, the diffuse contribution is negligible. Hence, when the surface is smooth enough to produce a well defined specular beam, the surface roughness can be determined from the relative intensity of the specular beam (Bennett and Porteus, 1961; Marx and Vorburger, 1990). Later, Porteus (Porteus, 1963) extended this study to shorter wavelengths and discussed the inadequacy of parameters such as the RMS roughness, the RMS slope and the correlation length for describing the reflectance for these wavelengths. In particular, Porteus (Porteus, 1963) showed how the distribution of heights of the surface irregularities can be determined from reflectance measurements at normal incidence.

2.2.4 Phenomenological Reflectance Models

A compromise between the empirical and physics-based methods is offered by the phenomenological approach adopted by Wolff (Wolff, 1994a), and, by Oren and Nayar (Oren and Nayar, 1995) which attempts to incorporate simple physics. These models attempt to account for departures from Lambertian reflectance due to surface roughness in a parametrically efficient manner. Each model modifies Lambert's cosine law in a different way. Oren and Nayar (Oren and Nayar, 1995) address the problem of rough reflectance from matte surfaces. This model relies on the assumptions that the surface is composed of V-groove microfacets and that the reflectance from each microfacet is Lambertian. In simplified form, their model accounts for the angular distribution of surface cavity walls by adding to the Lambert cosine law a term that depends on the squared sine of the incidence angle. This term results in brightening at the occluding object boundary. Wolff (Wolff, 1994a; Wolff, 1994b), on the other hand, has a physically deeper model for diffuse reflectance from shiny but slightly-rough surfaces. The model uses an angle

dependent Fresnel term to account for the refractive attenuation of incident light at the surface-air boundary. This Fresnel term modifies in a multiplicative way the Lambertian cosine model. The effect is to depress the surface radiance for near-normal incidence. Later, Wolff et al. (Wolff et al., 1998) suggested a combination of these two separate models for surfaces with intermediate roughness. The model includes both multiplication of the cosine law by a Fresnel term, and the addition of a sine-squared term. Hence, there can be both an increase in surface radiance at the limbs and a reduction in surface radiance at near-normal incidence.

van Ginneken et al. (van Ginneken et al., 1998) observe that a major drawback of the Oren-Nayar model is that it relies on the assumption of surface isotropy and this in turn requires that the cavities associated with roughness have no preferred direction. Such isotropic cavitated surfaces cannot exist, because they are inconsistent with surface geometry. Based on this observation, van Ginneken et al. have developed a model that can be used to predict reflectance from isotropic rough surfaces that have both diffuse and specular components (van Ginneken et al., 1998). The parameters of their model are the surface roughness measured in terms of the RMS slope, the surface albedo, and the balance between the diffuse and specular components. The model can also be used to predict changes in surface texture as the illumination and viewing angles vary. Although the model includes masking and shadowing effects using geometrical optics, the resulting integral for diffuse reflectance can not be evaluated analytically and must be calculated numerically.

2.3 Other Related Topics

As mentioned earlier, the problem of recovering shape from shading information has attracted researchers from different communities such as computer vision, computer graphics, optical physics, psychophysics, etc. Some of the most pertinent papers were reviewed in Sections 2.1 and 2.2. In this section we review some of the contributions which are related to the topic of the thesis in a more indirect way.

2.3.1 Multiple Light Sources

The majority of recent methods for non-Lambertian SFS work under multiple lighting conditions. The process of recovering shape from multiple shaded images is referred to as photometric stereo. Using multiple images not only improves the task of shape recovery, but also allows the estimation of surface albedo, surface reflectance properties, reflectance components and light-source direction to be rendered effectively tractable.

Traditional stereo is the estimation of 3D shape from two images taken by cameras which are slightly tilted along their optic axis (Cryer et al., 1995). Unfortunately, it is difficult to determine the correspondence between picture elements. Photometric stereo, on the other hand, varies the direction of incident illumination between successive images while holding the camera in a fixed position (Woodham, 1980). Hence, since the imaging geometry is not changed, there is no need to find the correspondence between image points. Shape from photometric stereo was first introduced by Woodham (Woodham, 1980) who showed that the surface gradients could be recovered by using only two input images of Lambertian surfaces. Later, he (Woodham, 1981) solved for both the surface normal and the albedo of Lambertian surfaces using three images with non-coplanar light-source directions.

For non-Lambertian shape recovery, Tagare and deFigueiredo (Tagare and deFigueiredo, 1991) proposed a theory of photometric stereo for a large class of non-Lambertian diffuse reflectance maps. They investigated whether three light-sources were sufficient to estimate the surface normal direction and the illumination strength. Ikeuchi (Ikeuchi, 1981) used the double-delta specular model to determine the shape of specular surfaces by photometric stereo. Nayar et al. (Nayar et al., 1990) developed the photometric sampling method using a hybrid reflectance model. The model used both specular and Lambertian components to extract the shape and reflectance of surfaces. Lee and Bajcsy (Lee and Bajcsy, 1992) proposed a method for detecting specularities using multiple color images from different directions. More recently, Solomon and Ikeuchi (Solomon and Ikeuchi, 1996) have described a four-source photometric stereo method with calibrated lighting

for determining albedo in the presence of specularity using the Torrance-Sparrow model (Torrance and Sparrow, 1967). Yuille and Snow (Yuille and Snow, 1997) showed how Hayakawa's approach (Hayakawa, 1994) can be generalized to include ambient background illumination and specularities. They also demonstrated the use of the integrability constraint to ensure that the set of surface normals forms a consistent surface.

Several methods recover a more accurate shape of the object as the number of photometric images increases. For instance, Lee and Kuo (Lee and Kuo, 1992) introduced parallel and cascade photometric stereo. The former takes all the photometric images together to produce the best estimation of the surface, while the latter takes the images, one-by-one. Later, Zhang et al. (Zhang et al., 1996; Zhang and Shah, 1997) have attempted to improve Lee and Kuo's approach to facilitate the task of shape recovery.

2.3.2 Shading and Illumination

The importance of illumination conditions for surface perception has been identified in psychophysics. For instance, Curran and Johnston (Curran and Johnston, 1996) have observed that it is the viewer's judgement about the direction of illumination which determines whether the surface curvature features in shaded scenes are perceived as either concave or convex. Langer and Bulthoff (Langer and Bulthoff, 2001) have presented psychophysical evidence which rejects the common belief that humans exhibit lighting invariant object recognition. They have examined the perception of qualitative local shape from shading and have found that significant variations occur as a function of the lighting conditions. Further, Fleming et al. (Fleming et al., 2001) have studied how humans perform the under-constrained task of estimating surface reflectance properties by just viewing single images and with no knowledge of the illumination conditions. Also, Dror et al. (Dror et al., 2001) have developed a machine vision system to perform reflectance estimation tasks automatically and in a manner similar to that performed by humans.

The importance of illumination in SFS has been widely emphasized in the literature. For instance, Zhang and Shah (Zhang and Shah, 1997) have noted that when only one

input image is used, in order to recover an accurate shape, most of the object should be illuminated. Belhumeur and Kriegman (Belhumeur and Kriegman, 1996) have tried to explain the reason why the same object can appear different when viewed in fixed pose.

There have been several attempts to estimate the illumination direction. In one of the earliest, Pentland (Pentland, 1982) described how the illuminant direction may be recovered for a region of an image by analyzing the distribution of changes in image intensity. Zheng and Chellappa (Zheng and Chellappa, 1991) have presented some techniques for determining the illuminant direction and surface albedo under Lambertian reflectance conditions. Finally, the Samaras and Metaxas method (Samaras and Metaxas, 1999) can be used to estimate the shape and the lighting direction using a two step iterative process.

2.3.3 Textured Surfaces

The analysis of rough surface reflectance is an important topic not only in shape-from-shading but also in shape-from-texture. However, although both processes are concerned with the recovery of local surface orientation, and hence height, they are based on very different computational models. In shape-from-texture, it is the texture gradients, measured either in terms of changes in texture primitive shape or texture primitive density, that is used to estimate local surface orientation. However, recent work on rough surface reflectance has blurred the distinction between shape-from-shading and shape-from-texture, which in many ways is an artificial one. Classical work on shape-from-texture was based on the assumption that the texture pattern is painted on the surface. While this may prove realistic for man-made regularly textured surfaces, for natural irregular textures it is unrealistic. A more realistic model is a photometric one which provides a physical account of the interaction of light with the structure of the surface. In other words a model of surface roughness is required. We have reviewed some of these models earlier in this chapter.

Fortunately, recent theories of surface texture have moved away from the naive idea that texture is painted onto the surface. Instead, attempts are being made to understand the formation of texture in terms of surface relief distributions. There are several concrete

examples of this approach. Leung and Malik (Leung and Malik, 2001) have introduced the idea of 3D textons. In a series of papers, Chantler and McGunnigle (Chantler and McGunnigle, 2000a) have studied the problem of texture classification. For instance, they have proposed a surface-rotation invariant scheme which uses the eigenvalues of gradient-space distribution of the surface as its features. They (Chantler and McGunnigle, 2000b) have also shown that when the illuminant direction is varied, the texture classification schemes fail due to the variable appearance of texture features. Moreover, Chantler et al. (Chantler et al., 2002) have recently looked critically at the problem of misclassification of surface textures due to changes in the angle of illumination incident. They have presented theoretical and experimental results that allow the behavior of texture features to be predicted as a function of illumination angle of tilt. This work may be utilized for illuminant tilt invariant classification. In an earlier paper, they (McGunnigle and Chantler, 2000) also developed a model of the angular dependence of texture appearance for rough surface reflectance under different illuminant directions. Although this approach has lead to good models of the appearance of rough surfaces, it has not lead to practical shape-from-texture algorithms.

BRDF Approach

Reflectance data can be obtained in different ways. An alternative to developing a detailed method is to acquire empirical estimates of the bidirectional reflectance function (BRDF) from data. The BRDF is defined to be the ratio of the outgoing surface radiance to the incoming light-source irradiance per unit solid angle. Measuring the BRDF is hence an expensive and time consuming process since it has four degrees of freedom corresponding to the zenith and azimuth angles of the light-source and the viewer. For instance, Dana et al. (Dana et al., 1997) have reported the results of a study in which the BRDF's of a number of different types of 3D surface textures have been catalogued. Later, Dana and Nayar (Dana and Nayar, 1998) demonstrated that their model can be used to account for the histogram of pixel intensities present in image texture. More recently, they (Dana and Nayar, 1999) presented a correlation model for a particular class of 3D texture due to

surface height variation. Finally, Magda et al. (Magda et al., 2001) have commented on how shape can be recovered from surfaces with arbitrary BRDF's.

Although there are clearly a number of potential reflectance models that can be used to describe reflectance from non-Lambertian surfaces, there has been relatively little effort aimed at comparing their predictions against empirical reflectance data. As mentioned earlier, one of the reasons for this is that performing measurements of the BRDF is a time consuming and laborious process. For instance, Watson and Raven (Watson and Raven, 2001) have compared the fit of various parameterized models, including the Oren-Nayar and the Wolff models, to experimental BRDF data from a variety of surfaces. The surfaces used in the study represent the extremes of material properties found in the real world.

There are alternative approaches to estimating the BRDF which are less time consuming. For instance, a method for estimating radiance against incidence angle from images of curved surfaces has been proposed by Robles-Kelly and Hancock (Robles-Kelly and Hancock, 2003). The method is simple and can be used in conjunction with SFS. Finally, Zickler et al. (Zickler et al., 2002) have reported a Helmholtz stereo method that can recover surface normal directions by reversing the position of the light-source and the camera. Their method makes no assumptions about the nature of the BRDF's of objects.

2.4 Conclusions

From the review of different reflectance models, the main points to conclude are as follows. The modelling of rough surface reflectance is important in both computer vision and computer graphics, and has been the subject of sustained research activity for some four decades. In fact, the quest for a reflectance model that can accurately account for the observed surface radiance under a variety of roughness conditions, and a variety of viewing geometries has proved to be an elusive one. Specifically, although specularities provide important geometric constraints for SFS, their reliable modelling, detection and separation in single-band images has proved an elusive goal.

From the literature review presented in this chapter we can also conclude that non-

Lambertian SFS is still an active and open research area. It is also clear that most of the recent non-Lambertian SFS methods apply only to photometric stereo and use more than one image input. Moreover, such methods do not employ reflectance models and use alternative methods such as reflectance ground-truth or the reflectance map. However, in this thesis we use just a single grey-scale image as input and show how we can bridge the gap between conventional SFS algorithms and non-Lambertian reflectance models

Because light reflection from most real-world surfaces does not follow Lambert's law, any progress made by Lambertian SFS algorithms towards recovering more accurate surface normal directions cannot be a major one. The reason for this is that such SFS algorithms ignore the effect of reflectance angle on the process of reflection and just assume that at every surface point the incidence angle is equal to the arc-cosine of the image brightness. As a result, when the light-source direction is aligned with the z axis, their main assumption is that the zenith angle of each surface normal is known and fixed, and their challenge is to recover the corresponding azimuth angles.

Nonetheless, there have been some attempts towards using a combination of specular and Lambertian reflectance models in conjunction with SFS. Although this approach results in the recovery of better surface shape, it cannot deliver fine results for surfaces whose diffuse reflectance component does not follow Lambert's law. Hence, we not only include a specular reflectance model in the process of SFS, but also make use of several non-Lambertian diffuse reflectance models to extract the corresponding Lambertian component. The corrected Lambertian image can be used as input to Lambertian SFS methods.

Finally, the importance of exploiting non-Lambertian reflectance models for SFS does not mean that we ignore the importance of the SFS method itself. Hence, we also present a new smoothing method which is an alternative to the existing smoothness and curvature consistency constraints.

Chapter 3

Highlight Removal using Specular Shape-from-Shading

The observation underpinning this chapter is that although considerable effort has gone into the recovery of accurate surface geometry, existing shape-from-shading (SFS) methods are confined to situations in which the reflectance is predominantly Lambertian. When local specularities are present, then the estimated geometry may be subject to error. The main problem that can occur is that specular highlights may be misidentified as curved surface features such as ridges and domes. Nonetheless, if specular highlights can be accurately located, then they can provide important cues that can be used to constrain the recovery of surface shape.

Our aim in this chapter is to develop a Bayesian framework for removing local specular highlights. We adopt a maximum *a posteriori* probability estimation framework using the iterative conditional modes algorithm of Besag (Besag, 1986). The aim is to estimate the proportions of Lambertian and specular reflectance. Once these proportions are to hand, then we can perform specularity removal. By applying a conventional SFS method to the resulting Lambertian reflectance image, we can recover more accurate surface normal directions in the presence of specularities.

To meet this goal we must provide probability distributions for the surface normals associated with the observed image intensity. Following Worthington and Hancock (Wor-

thington and Hancock, 1999), we develop a geometric framework and concentrate on modelling the angular distributions for the surface normals resulting from the two reflectance components. In the case of Lambertian reflectance, the surface normal is constrained to fall on an irradiance cone whose axis is aligned in the light-source direction and whose apex angle is determined by the observed image brightness. We assume that angular deviations from this irradiance cone follow a Gaussian distribution. For specular reflectance, the surface normal direction is the bisector of the viewer and light-source directions. We investigate three alternative models for the specular component. Here we use the Beckmann distribution (Beckmann and Spizzichino, 1963), and, the Gaussian simplification of Healey-Binford (Healey and Binford, 1987) to model the angular distribution of surface normals associated with specularities. As an alternative to modelling the directions of surface normals, we also investigate the use of the Torrance-Sparrow model (Torrance and Sparrow, 1967) to describe the distribution of specular intensities.

At each image location, the Bayes framework allows us to estimate the proportions of Lambertian and specular reflectance. With this information to hand we can compute the weighted mean of the ideal Lambertian and specular reflectance directions, and hence obtain the posterior-mean surface normal direction. Whereas, the Lambertian surface normal falls on the irradiance cone and the specular surface normal bisects the viewer and light-source directions, the posterior-mean surface normal is a weighted average of these two directions. In practice, we keep the directions of the specular surface normals fixed and only update the positions of Lambertian surface normals on their local irradiance cones. Hence, at each image location we need only estimate the polar angle of the Lambertian surface normal on the irradiance cone. Knowledge of this angle together with the probability distributions for Lambertian and specular reflectance, allows us to estimate the reflectance proportions. Hence, the iterated conditional modes algorithm alternates between interleaved steps for estimating posterior-mean surface normal directions and reflectance proportions.

Once the iterated conditional modes algorithm has converged, then from the available posterior-mean surface normal direction, we can reconstruct the specular reflectance us-

ing the Torrance-Sparrow model. This allows us to remove specular highlights from the original image. By re-applying the SFS technique to the residual matte reflectance component, we are able to extract better surface normal information which is free of the spurious curvature artifacts associated with specular highlights. Experiments on real world images reveal promising results. Ground-truth data is provided by imagery obtained with crossed polaroid filters. This reveals not only that the method accurately estimates the proportion of specular reflection, but that it also results in good surface normal reconstruction in the proximity of specular highlights.

The outline of the chapter is as follows. In Section 3.1 we discuss the geometry of the Lambertian and specular reflectance components. Section 3.2 describes the iterative conditional modes algorithm and its implementation. Sections 3.3 and 3.4 respectively address the probabilistic modelling of specular and Lambertian reflectance. In Section 3.6 we describe the Torrance-Sparrow model may be used for specularity removal. Experiments are described in Section 3.7. Finally, conclusions and directions for further work are presented in Section 3.8.

3.1 Reflectance Geometries and Models

In this section we outline the geometry of the reflectance processes which underpin our specularity subtraction model. We adopt a two-component model in which the predominantly Lambertian surface reflectance exhibits local specular highlights.

3.1.1 Specular Reflectance

The first component of our reflectance process is concerned with modelling local specular highlights on the observed surface. For specular reflection the surface normal, the light-source direction and the viewing direction are coplanar. The incidence angle (i.e. the angle between the surface normal and the light-source direction) is equal to the angle of specular reflectance (i.e. the angle between the surface normal and the viewing direction). Hence, for specular reflection, the direction of the surface normal $\vec{N}_S$ is the bisector of

the light-source direction $\vec{L}$ and the viewing direction $\vec{V}$ and the unit-vector is

$$\vec{N}_S = \frac{(\vec{L} + \vec{V})}{\|\vec{L} + \vec{V}\|} \tag{3.1}$$

It is important to stress that the surface normal for the specular reflectance component is fully constrained if the light-source and the viewing directions are known. We therefore keep the specular surface normal direction fixed throughout our iterative recovery of the needle map. However, for consistency in the iterative process, we use $\vec{N}_S^{(n)}$ to denote the specular surface normal direction. We also refer to this direction as specular spike direction (which is sometimes defined $\vec{H}$ in literature). The geometry of the specular reflectance process is illustrated in Figure 3.1.

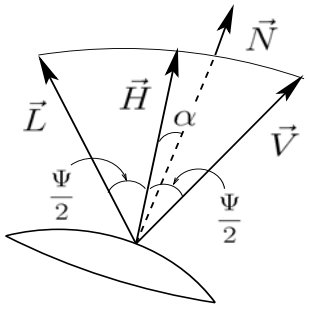

Figure 3.1: A simple geometry of the specular reflection process.

There are many locations on a shiny surface whose corresponding surface normal directions are unconstrained but deviate slightly from specular spike direction. However, the contribution of the reflected light from such locations in the observed specular intensity is considerable. In practice, the direction of specular reflectance is not restricted to the direction of the specular spike. According to the microfacet model, the surface is composed of microscopic planar facets that behave like mirrors. The specular intensity due to variations in surface microfacet orientation need to be computed using a model. Perhaps the most commonly used model for specular intensity is the Torrance-Sparrow model. Because we use this model both for reconstruction of specular intensity and for modelling the distribution of surface normal directions, here we explain the model and

describe the different parameters involved.

Torrance-Sparrow Model for Specular Intensity

The Torrance-Sparrow model assumes that the surface is composed of a collection of randomly oriented mirror-like microfacets for which Fresnel theory individually applies to each microfacet. Also, the specular component of the reflected light is assumed to come from those facets oriented in the direction of the specular spike $\vec{N}_S$. The specular intensity at the pixel with coordinates (i, j) and the mean surface normal $\vec{M}$ is given by

$$I_S(i,j) = \frac{fGD(\alpha)}{\pi(\vec{M}\cdot\vec{V})(\vec{M}\cdot\vec{L})} \tag{3.2}$$

The model is controlled by four terms. The first of these is the Fresnel reflection term f which is a function of incidence angle θ_i and the index of refraction n. Using trigonometry, Cook and Torrance (Cook and Torrance, 1981) simplified the Fresnel coefficient for near-normal incidence to

$$f = \frac{1}{2}\frac{(b-c)^2}{(b+c)^2}\left\{1 + \frac{(c(b+c)-1)^2}{(c(b-c)+1)^2}\right\} \tag{3.3}$$

where $c = \vec{V}\cdot\vec{N}_S$ and $b^2 = n^2 + c^2 - 1$. This form of the Fresnel coefficient is used here since it is relatively simple and also precise enough for small incidence angles. However, there is no reason why the general Fresnel coefficient cannot be used. This model is discussed in details by Angelopoulou and Poger (Angelopoulou and Poger, 2003). Moreover, another expression for the Fresnel coefficient used by Wolff et al. (Wolff et al., 1998) in modelling diffuse reflectance from smooth surfaces and is discussed later-on in details in Chapter 4. It is interesting to note that f has a substantially different form for metallic and nonmetallic materials (Blinn, 1977). For metals, f is close to unity for large values of n. For nonmetals and small values of n, it has a more exponential character varying from zero at $\theta_i = 0$ to unity at $\theta_i = \pi/2$. Here, we are only interested in nonmetals, and so, the approximation $f \approx 1$ is not applicable. For dielectrics, the index of refraction n is in the

range from 1.4 to 2.0. However, as noted by Wolff et al. (Wolff et al., 1998), f is weakly dependent upon the physical parameter n for a wide variety of dielectrics.

Secondly, there is the geometrical attenuation factor

$$G = min\left[1, 2\frac{(\vec{M}\cdot\vec{N_S})(\vec{M}\cdot\vec{V})}{(\vec{V}\cdot\vec{N_S})}, 2\frac{(\vec{M}\cdot\vec{N_S})(\vec{M}\cdot\vec{L})}{(\vec{L}\cdot\vec{N_S})}\right] \tag{3.4}$$

Here G represents the contribution of masking and shadowing of microfacets by adjacent microfacets. If a microfacet reflects the entire incident light, then $G = 1$. However, for a microfacet where masking effects are present, i.e. part of the reflected beam is blocked by a nearby microfacet, then $G = 2(\vec{M}\cdot\vec{N_S})(\vec{M}\cdot\vec{V})/(\vec{V}\cdot\vec{N_S})$. Finally, for a microfacet where shadowing effects occur, i.e. part of the incident beam is blocked by a nearby microfacet, then $G = 2(\vec{M}\cdot\vec{N_S})(\vec{M}\cdot\vec{L})/(\vec{L}\cdot\vec{N_S})$.

Thirdly, there is the facet slope distribution function $D(\alpha)$ which represents the distribution of microfacet normals around the specular spike direction. There are a variety of distribution functions discussed in the literature including the Gaussian distribution function (Blinn, 1977). We follow Cook and Torrance (Cook and Torrance, 1981) who used the Beckmann distribution function for rough surfaces and which is given by

$$D(\alpha) = \frac{1}{\sigma_m^2\cos^4(\alpha)}\exp\left[-\left(\frac{\tan(\alpha)}{\sigma_m}\right)^2\right] \tag{3.5}$$

where α is the angle between the mean surface normal $\vec{M}$ and the specular spike direction $\vec{N_S}$, i.e. $\alpha = \cos^{-1}(\vec{M}\cdot\vec{N_S})$. Also, σ_m is the RMS slope which determines the spread of the specular component. Small values of σ_m signify gentle facet slopes and produce a distribution that is highly collimated around the specular spike direction. Large values of σ_m imply steep facet slopes and produce a distribution that is highly dispersed. Despite requiring more computation, this function gives the absolute magnitude of the reflectance without introducing arbitrary constants (Cook and Torrance, 1981). The normalized form of this function is also used in Section 3.3.1.

Fourthly, and finally, the two inner products, i.e. $(\vec{M}\cdot\vec{L})$ and $(\vec{M}\cdot\vec{V})$, represent the

cosine of the incidence and reflectance angles, i.e. $\cos(\theta_i)$ and $\cos(\theta_r)$, respectively. Note that in the specular intensity model given by Equation (3.2), $D(\alpha)$ and G only describe the light pathways and their geometric interactions with the surface microstructure, while the Fresnel reflection term describes how light is reflected from each smooth microfacet (Angelopoulou and Poger, 2003).

3.1.2 Lambertian Reflectance

In the case of Lambertian reflectance from a matte surface of constant albedo illuminated by a single collimated light-source, the observed intensity is independent of the viewing direction. The observed intensity depends only on the quantity of absorbed light, and this in turn is proportional to the cosine of the incidence angle between the light-source direction and the local normal vector to the surface. Suppose that $\vec{L}$ is the unit-vector in the direction of the light-source and that $\vec{N}_L(i,j)$ is the unit-vector in the surface normal direction for Lambertian reflectance at the pixel (i,j). For a surface of constant albedo, according to Lambert's law, the normalized image intensity at the pixel (i,j) is

$$E(i,j) = \vec{N}_L(i,j) \cdot \vec{L} \tag{3.6}$$

Lambert's law provides insufficient information to uniquely determine the surface normal direction. However, as recently observed by Worthington and Hancock (Worthington and Hancock, 1999), the equation does have a simple geometric interpretation which can be used to constrain the direction of the surface normal. The equation specifies that the surface normal must fall on the surface of a cone (as shown in Figure 3.2) whose axis is aligned in the light-source direction $\vec{L}$ and whose apex angle is $\theta_i = \cos^{-1}(E)$.

Worthington and Hancock (Worthington and Hancock, 1999) exploit this property to develop a two-step iterative process for SFS. The process commences from a configuration in which the surface normals are placed on the position on the irradiance cone where their projections onto the image plane are aligned in the direction of the local (Canny) image gradient. This geometry is illustrated in Figure 3.2.

In our specular SFS scheme, where the specular surface normal is fixed in direction, the only free parameter at each image location is the rotation angle of the Lambertian surface normal on the irradiance cone about the light-source direction. This angle is found using a maximum-likelihood method to locate the rotation angle.

3.2 Bayesian Framework

The aim in this chapter is to develop a Bayes-decision scheme for estimating the reflectance proportions of the Lambertian and specular reflectance modes. In other words, we wish to compute the *a posteriori* probabilities of specular and Lambertian reflectance. To meet this goal we make use of the iterated conditional modes (ICM) algorithm of Besag (Besag, 1986). Although the method has notoriously poor global optimization properties, we use it here because it is simple and relatively efficient. In this chapter our primary aim is to investigate the feasibility of applying probabilistic methods to the SFS problem. The use of more sophisticated optimization methods such as simulated annealing, Markov chain Monte-Carlo or mean field annealing is a topic of future research.

ICM is an iterative process that interleaves two computational steps. The first of these involves the maximum *a posteriori* probability estimation of hidden observations or labels. In the second step these estimates are used to compute contextual probabilities for putative updates of the hidden observations.

In our application of the iterated conditional modes algorithm we treat the field of surface normals as hidden data. At individual pixel sites, the surface normals can arise from either a Lambertian or a specular reflectance process. As demonstrated in Section 3.1, these two reflectance modes are characterized by distinct geometric processes. Accordingly, we characterize the surface normals arising from the two reflectance processes using separate class-conditional probability distributions. These distributions model the physics of the Lambertian and specular reflectance processes, and are detailed in Sections 3.4 and 3.5 of this chapter. In addition, we compute the conditional or with-context probabilities for individual surface normals given the pattern of surrounding surface normals.

The with-context probabilities model the smoothness of the field of recovered surface normals. At each image location, we maintain the following pieces of information. Firstly, there are the *a posteriori* probabilities or proportions for the Lambertian and specular reflectance components. Secondly, there are the modal surface normal directions for the two reflectance distributions. Here the modal surface normal is defined as the surface normal with the maximum probability. In the case of specular reflectance, the modal direction is fixed to be the bisector of the light-source and viewing directions. In the case of the Lambertian process, the modal surface normal rotates around the irradiance cone. Hence, at each image location there is a single free parameter to be estimated, namely, the rotation angle of the Lambertian surface normal about the light-source direction.

With these ingredients, the iterated conditional modes algorithm proceeds as follows. At each image location, we compute the posterior-mean surface normal. This is found by using the Lambertian and specular reflectance proportions to compute a weighted average of the modal surface normal directions. To recover the maximum likelihood Lambertian surface normal direction, we project the posterior-mean surface normal onto the nearest location on the Lambertian irradiance cone. Using the probability distributions for the surface normal directions together with the with-context probabilities, we update the *a posteriori* probabilities for the Lambertian and specular reflectance components. This allows us to update the posterior-mean surface normal and hence iterate the algorithm.

3.2.1 Formal Details

The aim is to label pixels according to the reflectance mode from which they originated. The class identity for the pixel (i, j) at iteration n is denoted by $\omega_{i,j}^{(n)}$. The class-identity may be drawn from the set $\Omega = \{S, L\}$ where S is the specular reflectance label and L is the Lambertian reflectance label. For each image location, we maintain a specular surface normal and a Lambertian surface normal which satisfy the geometric constraints outlined in Section 3.1. At iteration n of the algorithm the currently available estimates of the Lambertian and specular surface normals at the pixel indexed (i, j) are respectively

$\vec{N}_L^{(n)}(i,j)$ and $\vec{N}_S^{(n)}(i,j)$. In the case of the specular component, the normal direction is in the direction of local specular reflection, i.e. the bisector of the light-source and viewer directions, and does not change with iteration number. In the case of Lambertian reflectance, the surface normal direction varies with iteration number, but is always projected to be positioned on the irradiance cone.

To develop our decision process, we require two probabilistic modelling ingredients. The first of these are separate probability density functions which can be used to represent the distributions of surface normals for the specular and Lambertian components. We evaluate these densities at the posterior-mean surface normal $\vec{M}^{(n)}(i,j)$ computed at iteration n. The reason for doing this is that the current values of the specular and Lambertian normals are guaranteed to satisfy the geometric constraints outlined in Section 3.1. As a result, they will be associated with vanishing angular error. Accordingly, we let

$$q_{i,j}^{(n)}(L) = p(E(i,j)|\vec{M}^{(n)}(i,j), \omega_{i,j}^{(n)} = L) \tag{3.7}$$

be the probability density function of the brightness given the posterior-mean surface normal under the Lambertian reflectance model. Similarly, we let

$$q_{i,j}^{(n)}(S) = p(E(i,j)|\vec{M}^{(n)}(i,j), \omega_{i,j}^{(n)} = S) \tag{3.8}$$

denote the same probability under the specular reflectance model.

The second probabilistic ingredient is a smoothness prior for the selected surface normal direction. This component of the model incorporates contextual information. Indexing the surface normals according to their pixel locations, suppose that

$$\Gamma_{i,j}^{(n)} = \{\vec{M}^{(n)}(k,l)|(k,l) \in G_{i,j}\} \tag{3.9}$$

is the set of posterior-mean surface normals in the neighborhood $G_{i,j}$ of the pixel (i,j). We let

$$P_{i,j}^{(n)}(L) = P(\vec{N}_L^{(n)}(i,j)|\Gamma_{i,j}^{(n)}) \tag{3.10}$$

be the conditional probability of the Lambertian surface normal at the location (i,j) given the field of surrounding posterior-mean surface normals. With these ingredients, then according to iterated conditional modes the updated surface normal should be the one which has maximum *a posteriori* probability, given the observed image brightness and the current set of neighboring surface normals. In other words it satisfies the condition

$$\vec{M}^{(n+1)}(i,j) = \arg\max_{\vec{\hat{M}}(i,j)} P(\vec{\hat{M}}(i,j)|E(i,j),\Gamma_{i,j}^{(n)}) \tag{3.11}$$

To develop a decision scheme, we note that the observed brightness is sampled from a mixture of Lambertian and specular components. We make this mixture model explicit by writing

$$P(\vec{\hat{M}}(i,j)|E(i,j),\Gamma_{i,j}^{(n)}) = \sum_{\Lambda\in\Omega} P(\vec{\hat{M}}(i,j),\omega_{i,j}^{(n)} = \Lambda|E(i,j),\Gamma_{i,j}^{(n)}) \tag{3.12}$$

where Λ may be drawn from the set $\Omega = \{S, L\}$. Using the Bayes formula we have

$$\begin{aligned} P(\vec{\hat{M}}(i,j),\omega_{i,j}^{(n)} = \Lambda|E(i,j),\Gamma_{i,j}^{(n)}) &= \frac{1}{P(E(i,j),\Gamma_{i,j}^{(n)})} \\ \{P(E(i,j)|\vec{\hat{M}}(i,j),\omega_{i,j}^{(n)} = \Lambda,\Gamma_{i,j}^{(n)})P(\vec{\hat{M}}(i,j),\omega_{i,j}^{(n)} = \Lambda|\Gamma_{i,j}^{(n)})P(\Gamma_{i,j}^{(n)})\} \end{aligned} \tag{3.13}$$

As a result the *a posteriori* probability of the updated surface normal direction is

$$\begin{aligned} P(\vec{\hat{M}}(i,j)|E(i,j),\Gamma_{i,j}^{(n)}) &= \frac{P(\Gamma_{i,j}^{(n)})}{P(E(i,j),\Gamma_{i,j}^{(n)})} \\ \sum_{\Lambda\in\Omega}\{P(E(i,j)|\vec{\hat{M}}(i,j),\Gamma_{i,j}^{(n)},\omega_{i,j}^{(n)} = \Lambda)P(\vec{\hat{M}}(i,j),\omega_{i,j}^{(n)} = \Lambda|\Gamma_{i,j}^{(n)})\} \end{aligned} \tag{3.14}$$

Note that because the probability $P(E(i,j)|\vec{\hat{M}}(i,j),\Gamma_{i,j}^{(n)},\omega_{i,j}^{(n)} = \Lambda)$ only depends on the surface normal for object (i,j) and not on the surrounding field, it is equivalent to $P(E(i,j)|\vec{\hat{M}}(i,j),\omega_{i,j}^{(n)} = \Lambda)$. We can also replace $P(\vec{\hat{M}}(i,j),\omega_{i,j}^{(n)} = \Lambda|\Gamma_{i,j}^{(n)})$ by $P(\vec{N}_{\Lambda}^{(n)}(i,j)|\Gamma_{i,j}^{(n)})$ since $\vec{N}_{\Lambda}$ is the surface normal corresponding to the Lambertian or

specular modes at the pixel indexed (i,j). As a result we have

$$P(\vec{\hat{M}}(i,j)|E(i,j),\Gamma_{i,j}^{(n)}) = \frac{P(\Gamma_{i,j}^{(n)})}{P(E(i,j),\Gamma_{i,j}^{(n)})}$$
$$\sum_{\Lambda\in\Omega}\{P(E(i,j)|\vec{\hat{M}}(i,j),\omega_{i,j}^{(n)}=\Lambda)P(\vec{N}_{\Lambda}^{(n)}(i,j)|\Gamma_{i,j}^{(n)})\} \tag{3.15}$$

Finally, the probabilities $P(\Gamma_{i,j}^{(n)})$ and $P(E(i,j),\Gamma_{i,j}^{(n)})$ depend only on observation of fixed measurements and not on labels which are changing. Hence these probabilities play no role in our decision scheme. We select the surface normal direction which is the posterior-mean of the mixture of Lambertian and specular components, i.e.

$$\vec{M}^{(n+1)}(i,j) = \arg\max_{\vec{\hat{M}}(i,j)} \sum_{\Lambda\in\Omega}\{P(E(i,j)|\vec{\hat{M}}(i,j),\omega_{i,j}^{(n)}=\Lambda)P(\vec{N}_{\Lambda}^{(n)}(i,j)|\Gamma_{i,j}^{(n)})\} \tag{3.16}$$

To compute the posterior-mean, we require the *a posteriori* probabilities for the two reflectance modes. The *a posteriori* probability that the pixel (i,j) belongs to the Lambertian class at iteration n is

$$P(\omega_{i,j}^{(n)}=L|E(i,j),\vec{M}^{(n)}(i,j)) =$$
$$\frac{P(E(i,j)|\vec{M}(i,j),\omega_{i,j}^{(n)}=L)P(\vec{N}_{L}^{(n)}(i,j)|\Gamma_{i,j}^{(n)})}{\sum_{\Lambda\in\Omega}\{P(E(i,j)|\vec{M}(i,j),\omega_{i,j}^{(n)}=\Lambda)P(\vec{N}_{\Lambda}^{(n)}(i,j)|\Gamma_{i,j}^{(n)})\}} \tag{3.17}$$

Obviously, the *a posteriori* probability that the surface normal belongs to the specular class is the complement, i.e.

$$P(\omega_{i,j}^{(n)}=S|E(i,j),\vec{M}^{(n)}(i,j)) = 1 - P(\omega_{i,j}^{(n)}=L|E(i,j),\vec{M}^{(n)}(i,j)) \tag{3.18}$$

These two *a posteriori* probabilities can be used to represent the proportions of specular and Lambertian reflectance. With the *a posteriori* class probabilities to hand, we can update the estimate of the posterior-mean surface normal in the following manner

$$\vec{M}^{(n+1)}(i,j) = \sum_{\Lambda\in\Omega}\{\vec{N}_{\Lambda}^{(n)}(i,j)P(\omega_{i,j}^{(n)}=\Lambda|E(i,j),\vec{M}^{(n)}(i,j))\} \tag{3.19}$$

With the posterior-mean surface normal to hand, we can revise the direction of the modal Lambertian surface normal. To do this we project the smoothed posterior-mean surface normal onto the irradiance cone to give the updated Lambertian surface normal direction $\vec{N}_L^{(n)}(i,j)$. To apply this probabilistic framework, we require models for the probability density functions for the distributions of specular and Lambertian surface normal directions, and also for the conditional priors for the surface normals. In the subsequent sections of this chapter, we detail these models.

3.2.2 Algorithmic Steps

Having described the Bayes framework and the associated two-mode reflectance model, we are now in a position to develop a practical SFS algorithm. At each image location, we maintain a specular surface normal and a Lambertian surface normal which both satisfy the geometric constraints outlined in Section 3.1. The steps of the algorithm are summarized in Figure 3.2.

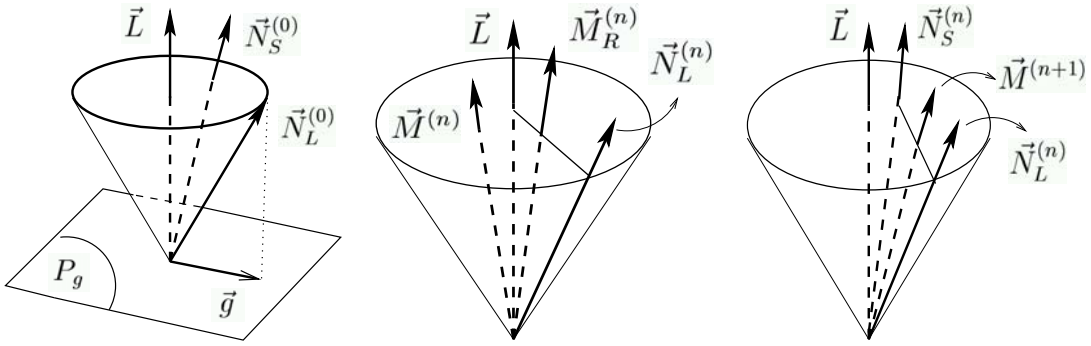

Figure 3.2: Geometry of the surface normal updating process: a) initialization of Lambertian surface normal $\vec{N}_L^{(0)}$ (initial posterior-mean surface normal $\vec{M}^{(0)}$ is identical to $\vec{N}_L^{(0)}$); b) smoothing the current posterior-mean surface normal $\vec{M}^{(n)}$ to $\vec{M}_R^{(n)}$, and, computing $\vec{N}_L^{(n)}$ by rotating $\vec{M}_R^{(n)}$ onto the cone ($\vec{N}_S^{(n)}$ is fixed to $\vec{N}_S$); c) computing the updated posterior-mean surface normal $\vec{M}^{(n+1)}$ using the current Lambertian and specular surface normals ($\vec{N}_L^{(n)}$ and $\vec{N}_S^{(n)}$) and the corresponding *a posteriori* probabilities.

We commence by initializing the algorithm. The initial Lambertian surface normal $\vec{N}_L^{(0)}(i,j)$ is constrained to lay on the irradiance cone in the direction of the image gradient. The specular surface normal is the bisector of the light-source and viewing directions.

The subsequent iterative steps of the algorithm (for iterations $n = 1, 2, ..$) are as follows:

- 1: The field of posterior-mean surface normals $\vec{M}^{(n)}(i,j)$ (initially identical to the field of Lambertian surface normals) is subjected to local smoothing. Here we use the curvature sensitive smoothing method of Worthington and Hancock (Worthington and Hancock, 1999). The smoothed surface normal is denoted by $\vec{M}_R^{(n)}(i,j)$.

- 2: We update the current estimate of the Lambertian surface normal direction by projecting the smoothed posterior-mean surface normal direction onto the nearest location on the irradiance cone using the geometric construction outlined in Section 3.1.2. This gives us the revised surface normal $\vec{N}_L^{(n)}$. Note that the specular surface normal direction $\vec{N}_S^{(n)}$ is fully constrained and hence does not need to be updated.

- 3: With the smoothed posterior-surface normal direction to hand, we compute the conditional measurement densities $q_{i,j}^{(n)}(L)$ and $q_{i,j}^{(n)}(S)$ for the Lambertian and specular reflectance modes. Taking the smoothed posterior-mean surface normal directions and the current surface normal directions for the two reflectance modes, we compute the smoothness priors $P_{i,j}^{(n)}(L)$ and $P_{i,j}^{(n)}(S)$. With these two ingredients in hand, we compute the updated with-context *a posteriori* probabilities for the specular and Lambertian reflectance modes.

- 4: Using the currently available estimates of the specular and Lambertian surface normals and the corresponding *a posteriori* probabilities, we compute the new posterior-mean surface normal $\vec{M}^{(n+1)}(i,j)$ and we return to step 1.

3.3 Distribution of Specular Component

As mentioned in Chapter 2, the modelling of specular reflectance is important for computer vision applications. Here we are interested in two approaches to the problem. The first of these is concerned with modelling specular intensities. Here the model of Torrance-Sparrow (Cook and Torrance, 1981; Torrance and Sparrow, 1967) captures the

physics of scattering by the microfacet structure of a surface. The second approach is to model the angular distribution associated with the reflected light in the proximity of specularities. Here the Beckmann distribution (Beckmann and Spizzichino, 1963) provides a relatively simple model which captures the angular shape of the specular spike. Healey and Binford (Healey and Binford, 1987) have an alternative model which can be used to model the distribution of specular intensities for regions of high surface curvature.

There are a number of relative advantages and disadvantages with the three models. Of the two angular models (the Beckmann, and, Healey-Binford), the Healey-Binford is the easier to control. The reason for this is that the Gaussian distribution underpinning the Healey-Binford model has only a single parameter (the width of the distribution) and maintains the same shape for all parameter settings. Although the Beckmann distribution also has only a single parameter, it is more difficult to control. The reason for this is that the distribution changes shape with different parameter settings. In particular, it varies between a single specular spike to two specular side lobes. In practice this proves difficult to control. In the case of the Torrance-Sparrow model which models the specular intensity, the main difficulty stems from the fact that the camera response is frequently saturated in the proximity of specularities. Hence, the model may not accurately account for the extreme intensities in the specularity, and this in turn may lead to poor shape estimation.

Bearing in mind the potential difficulties, in this section, we compare three different strategies for modelling the specular reflectance component. In doing this we explore the two different modelling approaches described above. Two of our models are based on the angular distribution of surface normal directions for specular reflection. Here we investigate the Beckmann distribution and a Gaussian approximation to it. The remaining model is based on the distribution of specular intensities using the Torrance-Sparrow model.

3.3.1 Beckmann Distribution

Here we are only interested in the distribution of surface normal directions for local specular highlights. Hence we use the Beckmann distribution (Beckmann and Spizzichino,

1963; Cook and Torrance, 1981) to model the angle $\alpha = \cos^{-1}(\vec{M}^{(n)}(i,j) \cdot \vec{N}_S^{(n)}(i,j))$ between the posterior-mean surface normal $\vec{M}^{(n)}(i,j)$ and the predicted direction of the specular spike $\vec{N}_S^{(n)}$. The normalized distribution is

$$q_{i,j}^{(n)}(S) = \frac{1}{\sigma_m(2+\sigma_m^2)\cos^4(\alpha)} \exp\left[-\left(\frac{\tan(\alpha)}{\sigma_m}\right)^2\right] \tag{3.20}$$

where σ_m is the RMS slope parameter which controls the angular shape of the distribution. This distribution can be used to model the shape of both the specular spike and the specular lobe. It makes no attempt to model the distribution of specular intensities.

The Beckmann distribution focuses on the physical modelling of the observed set of surface normal directions associated with specular highlights. It does not attempt to incorporate statistical uncertainties associated with poor surface normal reconstruction. The modelling of such statistical uncertainties is outside the scope of this chapter. One way to capture such uncertainties would be to use a Von Mises distribution (Mardia, 1999). The distribution which models both physical variability and statistical uncertainties would be a convolution of the Von Mises and Beckmann distributions. However, the resulting distribution is not tractable in closed-form.

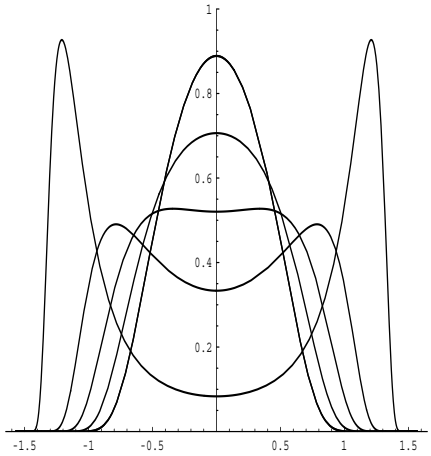

Figure 3.3: Dependence of the Beckmann distribution on the parameter σ_m.

3.3.2 Gaussian Distribution

One of the problems with the Beckmann distribution (Cook and Torrance, 1981) is that the shape is critically dependent on the width parameter σ_m. When σ_m is small, then the distribution is dominated by a peak close to $\alpha = 0$. When σ_m is large then there are two peaks at $|\alpha| = \pi/2$. This behavior is illustrated in Figure 3.3. The former case corresponds to a specular spike, while the latter case corresponds to a specular lobe.

Here we are mainly interested in modelling the angular shape of the specular spike. To simplify the control of the angular distribution, we have followed Healey and Binford (Healey and Binford, 1987) and have used a small angle approximation to the Beckmann distribution. The distribution is

$$q_{i,j}^{(n)}(S) = \frac{1}{\sqrt{2\pi}\sigma_\alpha} \exp\left[-\frac{1}{2}\left(\frac{\alpha}{\sigma_\alpha}\right)^2\right] \tag{3.21}$$

In other words the angle between the posterior-mean surface normal $\vec{M}^{(n)}(i,j)$ and the predicted direction of the specular spike $\vec{N}_S^{(n)}(i,j)$ follows a Gaussian distribution with zero mean and variance σ_α^2. It must be stressed that the distribution simply models deviations of the posterior-mean surface normal away from the direction of the specular spike and does not model the shape of the specular lobe. Again, it must be stressed that the modelling is physical rather than statistical. The relation between RMS slope σ_m in the Beckmann distribution function and slope angle σ_α in the Gaussian distribution function may be given by $\sigma_\alpha = \tan^{-1}(\sqrt{2}\sigma_m)$.

3.3.3 Intensity Distribution

Using the Torrance-Sparrow model for specular intensity (outlined in Section 3.1.1), we use an intensity-based distribution function as an alternative to the distribution of surface normal directions. Here, we assume that the observed specular intensities follow a

Gaussian distribution with variance σ_S^2. Under these assumptions we can write

$$q_{i,j}^{(n)}(S) = \frac{1}{\sqrt{2\pi}\sigma_S} \exp\left[-\frac{1}{2}\left(\frac{E(i,j) - I_S^{(n)}(i,j)}{\sigma_S}\right)^2\right] \tag{3.22}$$

This distribution models the variation in specular intensity resulting from the physical variation in specular reflectance direction. To exhibit the iterative nature of the process, here the specular intensity at iteration n is referred to as $I_S^{(n)}$. Also, the iteration dependant vectors $\vec{M}^{(n)}$ and $\vec{N}_S^{(n)}$ are used instead of $\vec{M}$ and $\vec{N}_S$ appearing in Equations (3.1) to (3.5).

3.4 Distribution of Lambertian Component

Our model of the Lambertian reflectance process assumes that the observed intensity values follow a Gaussian distribution with variance σ_L^2. The mean intensity is $(\vec{M}^{(n)} \cdot \vec{L})$. Under these assumptions we can write

$$q_{i,j}^{(n)}(L) = \frac{1}{\sqrt{2\pi}\sigma_L} \exp\left[-\frac{1}{2}\left(\frac{E(i,j) - \vec{M}^{(n)}(i,j) \cdot \vec{L})}{\sigma_L}\right)^2\right] \tag{3.23}$$

It must be stressed that there are alternative models to hand. For instance, Wolff et al. (Wolff et al., 1998) have looked critically at the validity of Lambert's law and have found departures at small grazing angles to the surface. These models require the specification of additional parameters which will be investigated in more detail in Chapter 4. Moreover, there are alternatives to the modelling of error using the difference between measured and predicted image brightness. For instance, we could use a Von Mises distribution to angular deviations of the surface normal from the Lambertian irradiance cone (Mardia, 1999).

3.5 Smoothness Prior

Our model for the surface normal smoothness prior is based on the average value of the inner product of the surface normal at the location (i, j) with the surrounding field of

posterior-mean surface normals. We write

$$P_{i,j}^{(n)}(\Lambda) = \frac{1}{2|G_{i,j}|}\left[|G_{i,j}| + \sum_{(k,l)\in G_{i,j}} \vec{N}_{\Lambda}^{(n)}(i,j)\cdot\vec{M}^{(n)}(k,l)\right] \tag{3.24}$$

When the posterior-mean surface normals from the neighborhood $G_{i,j}$ are aligned in the direction of $\vec{N}_{\Lambda}^{(n)}(i,j)$, then $P_{i,j}^{(n)}(\Lambda) = 1$. The larger the misalignment then the smaller the value of smoothness prior.

It must be stressed that there are alternatives to the above smoothness prior. For instance, Gauss-Markov models and Potts models have been widely used for the modelling of vector fields. Robust statistics have been used in the work of Black and Rangarajan (Black and Rangarajan, 1996) on motion segmentation. Worthington and Hancock (Worthington and Hancock, 1999) have developed a topography sensitive surface normal smoothing method which uses weighting functions suggested by robust statistics literature. However, here our aim is to use a simple method which does not require additional parameter control.

3.6 Specularity Removal

In this section we describe how the posterior-mean surface normals delivered by our SFS algorithm can be used for the purposes of reconstructing the specular intensity component. The reason for doing this is that the specular intensity may be removed from the original image intensity to give a corrected Lambertian image. By re-applying the SFS algorithm to this corrected intensity image, we aim to recover improved surface normal estimates, free of the high curvature artifacts of specular highlights. The process of removing the specular component can be viewed as equivalent to that of adjusting the angles of the Lambertian irradiance cones (Section 3.1.1). In principle, this specularity removal process could be iterated several times. However, in practice we find that one or at most two applications is sufficient to give insignificant specular probabilities over the entire image.

The simplest way to compute the Lambertian intensity $I_L(i,j)$ is to subtract the specu-

lar intensity $I_S(i,j)$ computed using the Torrance-Sparrow model (Torrance and Sparrow, 1967) from the measured image brightness $E(i,j)$. Obviously, the balance between specular and Lambertian reflection must be taken into account. In other words, at the pixel indexed (i,j) with the posterior-mean surface normal $\vec{M}(i,j)$, if the *a posteriori* probabilities are used to represent this balance, we can write

$$E(i,j) = \sum_{\Lambda \in \Omega} \{P(\omega_{i,j} = \Lambda | E(i,j), \vec{M}(i,j)) I_\Lambda(i,j)\} \tag{3.25}$$

Subsequently, the Lambertian intensity resulting from specular subtraction is

$$I_L(i,j) = \frac{E(i,j) - P(\omega_{i,j} = S | E(i,j), \vec{M}(i,j)) I_S(i,j)}{P(\omega_{i,j} = L | E(i,j), \vec{M}(i,j))} \tag{3.26}$$

Unfortunately, this method proves unstable and unreliable since the image intensity is frequently saturated in the proximity of specularities. For this reason, we have chosen to use the *a posteriori* probabilities of the Lambertian reflectance component to compute a weighted intensity average. For each pixel (i,j) we compute a weighted average of the image intensity over a set of neighboring pixels $T_{i,j}$. This neighborhood is selected to be such that the likelihood of specular saturation is small. The weighted Lambertian intensity is given by

$$I_L(i,j) = \frac{\sum_{(k,l) \in T_{i,j}} \{P(\omega_{k,l} = L | E(k,l), \vec{M}(k,l)) E(k,l)\}}{\sum_{(k,l) \in T_{i,j}} P(\omega_{k,l} = L | E(k,l), \vec{M}(k,l))} \tag{3.27}$$

There are many ways in which to select the set of pixels $T_{i,j}$, but here we proceed as follows. First, we construct level-sets for the specular intensity I_S computed using the Torrance-Sparrow model. These level-sets form closed contours on the image plane. Next, we locate the contours which bound the pixel regions for which the *a posteriori* probability $P(\omega_{i,j} = S | E(i,j), \vec{M}(i,j)) > Q_S$. The quantity Q_S is a specular probability threshold which is usually set to the value $Q_S = 0.1$. This ensures that the enclosed regions contain most of the significant specular structure and all of the locations of likely saturation in the image. For the pixel (i,j) which falls within one of these enclosed re-

gions, the set $T_{i,j}$ consists of the four pixels on the bounding contour which belong to the row i or the column j. Note that the specularity removal process is controlled by the single roughness parameter of the surface (e.g. RMS slope σ_m in Equation 3.5).

Once the corrected Lambertian intensity $I_L(i,j)$ is to hand, then we can use it as input to a re-application of our SFS algorithm. However, rather than initializing the algorithm with surface normals defined by the image gradient, we use the posterior-mean surface normals available from the application of the SFS algorithm to the raw image intensity.

3.7 Experiments

The images used in our experiments have been captured using an Olympus 10E camera. The objects studied are made of white porcelain and are hence shiny. Each object has been imaged under controlled lighting conditions in a darkroom. The objects have been illuminated using a single collimated tungsten light-source. The light-source direction is recorded at the time the images are captured.

To ground-truth the specularity removal method, we have used a pair of polaroid filters. Since specularities are the result of coherent scattering from surface microfacets, they can be extinguished by polarizing the light beam and observing the object from a crossed polaroid. We have placed the first filter between the light-source and the illuminated object. The second filter was placed between the illuminated object and the camera. For each object we have collected a pair of images. The first of these is captured when the second filter (between the camera and the object) is rotated until there is maximal extinction of the observed specularities. The second image is obtained when the polaroid is rotated through 90° , i.e. there is minimal extinction of the specularities. We refer to the polarization conditions of the former image as “crossed” and of the latter as “uncrossed”.

3.7.1 Examples

In Figures 3.4-3.6 we show the results obtained for three of the objects used in our study. The objects are a porcelain bear, a porcelain vase and a porcelain urn. In each figure the

order of the images is as follows. In the top row from left to right we show the image obtained with uncrossed polaroids, the image obtained with crossed polaroids and the difference between the two images. In the remaining four rows, the left-hand panel is the result obtained using the Beckmann model, the centre panel is the result obtained using the Gaussian model of Healey-Binford, and the right-hand panel is the result obtained using the Torrance-Sparrow model. In the second row we show the matte images obtained by specularity removal. The specular image component is shown in the third row of the figures. In the fourth row we show the needle maps obtained from the matte images. Finally, the fifth row shows curvedness estimated from the needle maps (i.e. the result of adding the two principal curvatures in quadrature).

There are a number of features in the different panels of the figures that deserve comment. From the uncrossed polaroid images it is clear that there are several quite small specular reflections across the surface of the bear. The vase has larger specularities on the neck and the centre of the bulb. The urn has a complex pattern of specularities around the left-handle. From the crossed polaroid images it is clear that most of the specular structure is removed. The difference images reveal the locations of the specularities. They also reveal differences at the limbs of the three objects. This is due to coherent scattering from microfacets when the light grazes the surface. This is a well-known departure from Lambertian scattering which is due to both surface roughness effects and subsurface scattering effects. These effects will be discussed in Chapters 4 and 5.

Turning our attention to the matte and specular images, it is clear that each of the three models gives good results. In each case the specular structure is removed and the matte appearance is close to that obtained with the crossed polaroids. Also, the pattern of specularities obtained in each case corresponds to that obtained by subtracting the crossed polaroid images.

We now turn our attention to the quality of the shape information returned by the SFS algorithm and consider the needle maps and the curvedness images. Here residual specularities manifest themselves as high curvature local surface features. No such features are evident in the images. Moreover, the surface shape and the fine details on the surfaces

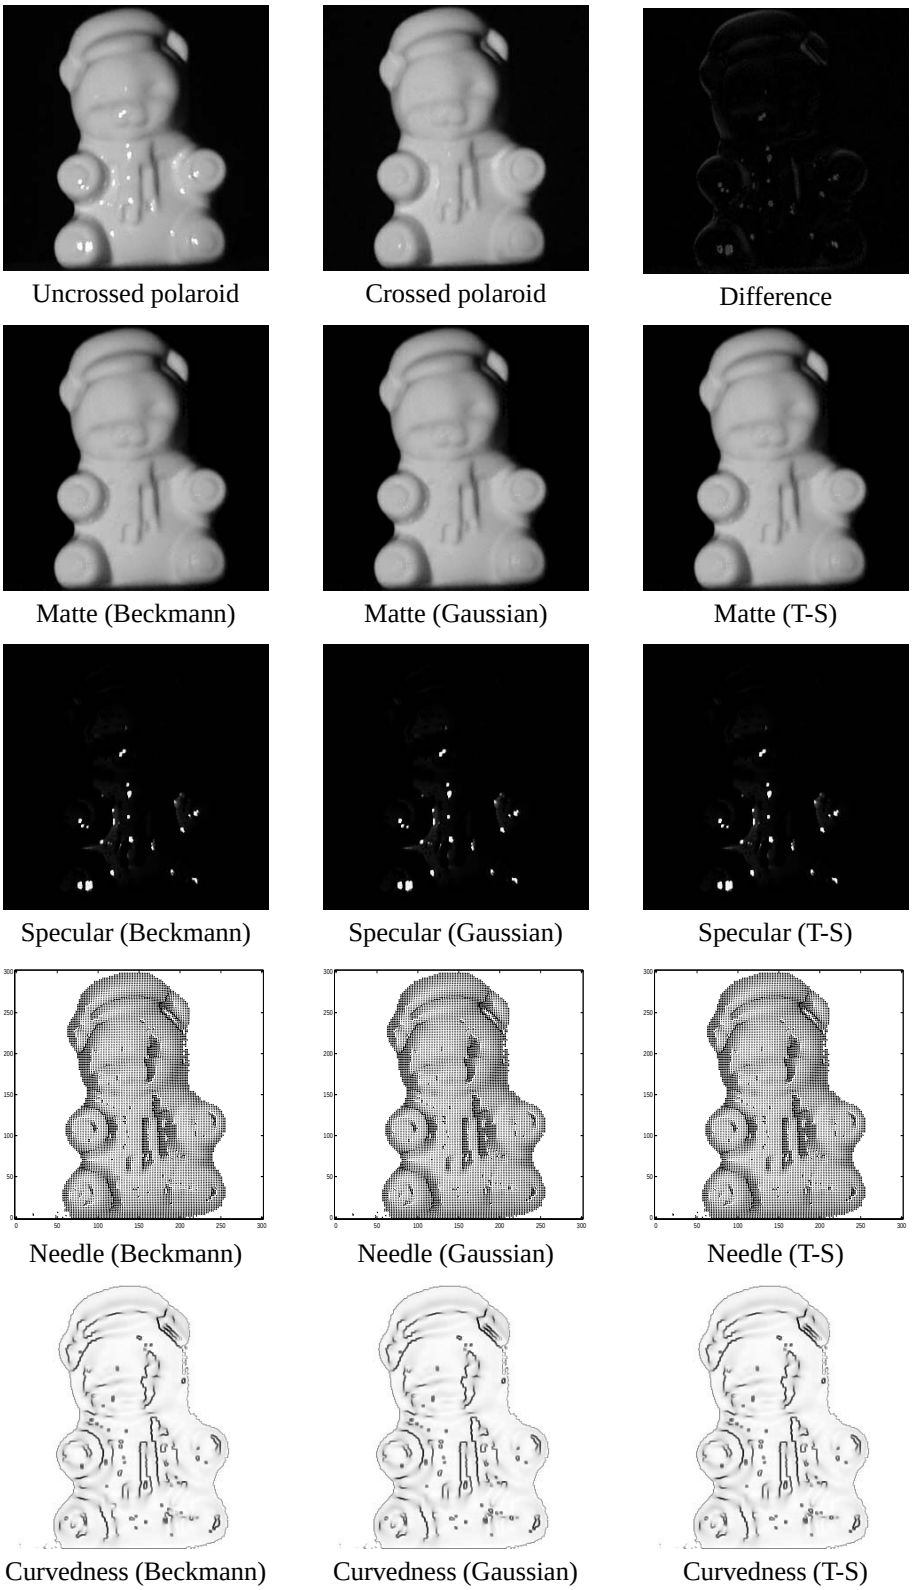

Figure 3.4: Highlight removal and shape recovery for a porcelain bear.

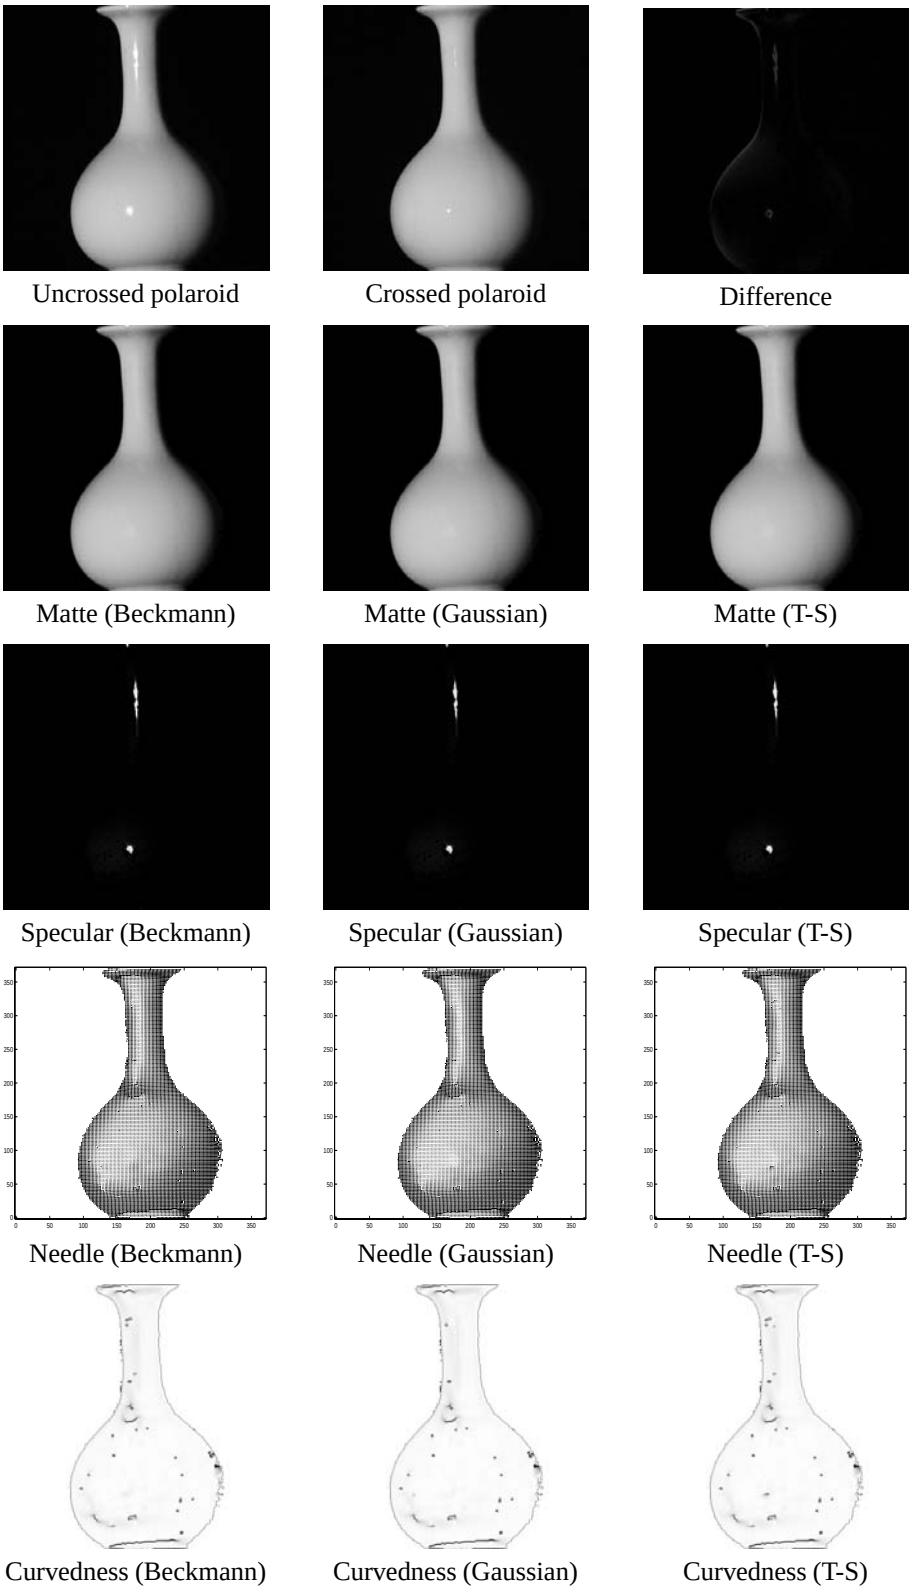

Figure 3.5: Highlight removal and shape recovery for a porcelain vase.

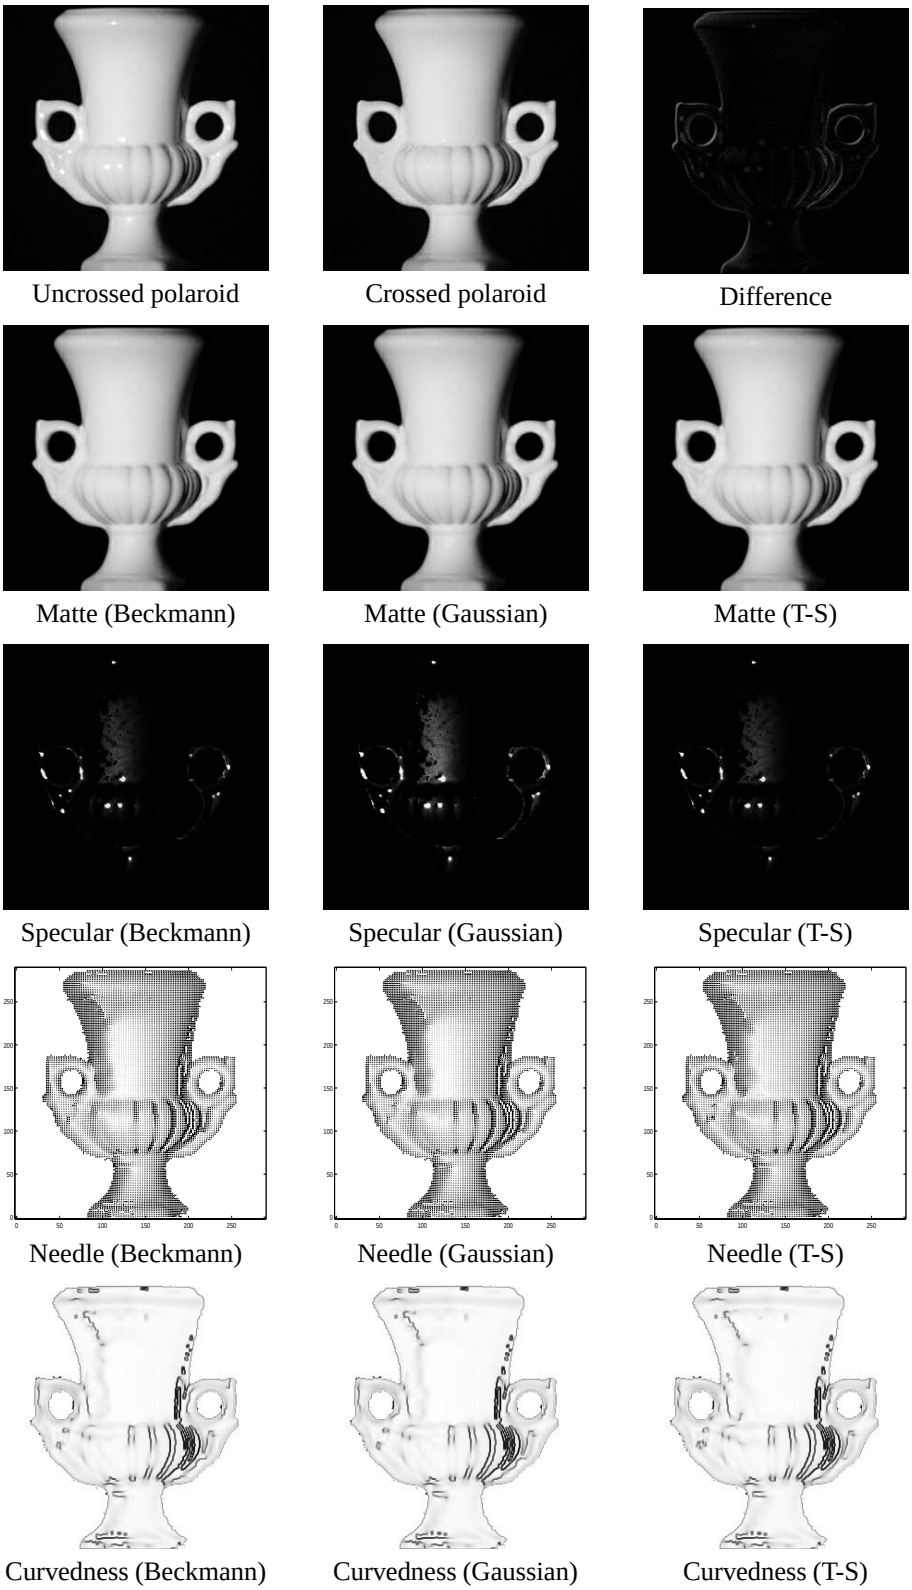

Figure 3.6: Highlight removal and shape recovery for a porcelain urn.

are well reconstructed. This is particularly evident from the curvedness image for the urn, where the ribbed structure is well detected. From these results it is also clear that there are no significant differences when the three specular models are used.

3.7.2 Comparison

In our next set of experiments, we compare the shape information extracted when the SFS method of Worthington and Hancock is applied to the uncrossed polaroid images and to the matte images obtained by specularity removal using the Beckmann distribution. In Figures 3.7-3.9 the panels are ordered as follows. In the top row the left-hand panel shows the needle map computed from the matte images obtained by specularity removal. The centre panel shows the needle map obtained from the raw uncrossed polaroid images. The right-hand panel shows the difference in the needle map directions. In the bottom row, the first panel shows the curvedness obtained from the matte images while the second panel shows the curvedness obtained from the uncrossed polaroid images.

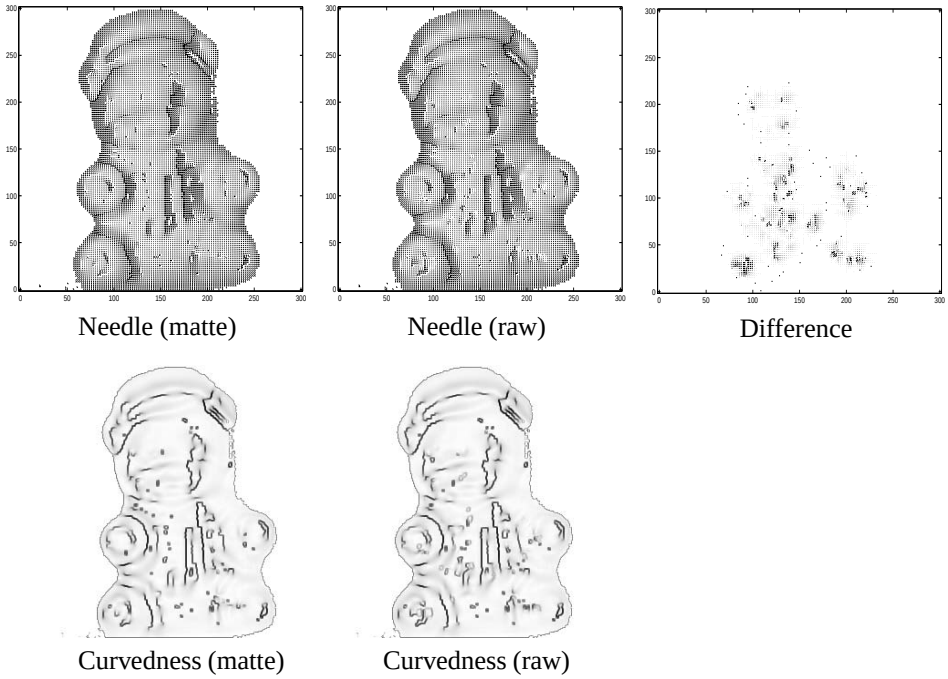

Figure 3.7: Comparing results corresponding to the raw and separated matte images.

The main features to note from these images are as follows. First, in the case of the

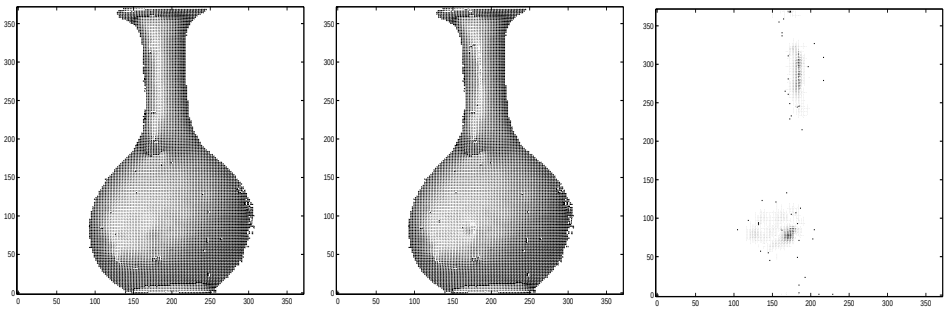

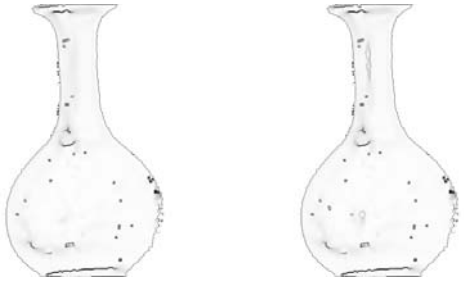

Figure 3.8: Comparing results corresponding to the raw and separated matte images.

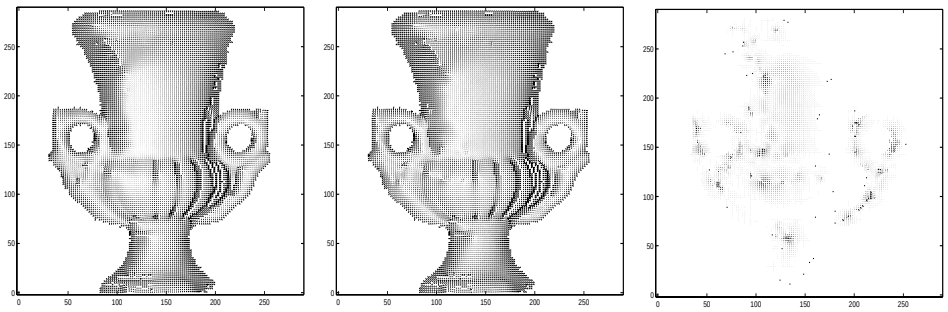

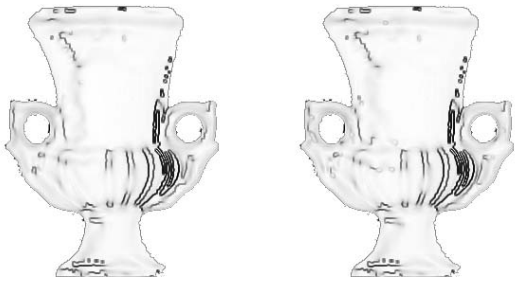

Figure 3.9: Comparing results corresponding to the raw and separated matte images.

raw images, there is significant local needle map structure around the specular highlights. In the case of our new specularity removal method, the structure is absent. From the difference-field for the two needle maps it is clear that the main structural differences also

occur around the specular highlights on the objects. These features are also evident in the curvedness images.

3.7.3 Specular Detail

To further illustrate the behavior of our method in the proximity of highlights, in Figures 3.11 and 3.12 we show some details for the porcelain urn shown in Figure 3.10. In Figure 3.11, we show a specular patch from the centre of the object while Figure 3.12 shows a region of quite complex specular reflection around the left-hand handle of the urn. In the two figures, the left-hand panel is the original detail, the centre panel shows the specular component, while the right-hand panel shows the matte component.

In Figure 3.13 we show Lambertian re-illuminations of the needle map extracted from the handle detail when the light-source direction is re-positioned. Here we move the direction of lighting from -40° to $+40^\circ$ to the image plane normal, in the direction of the x axis. The images reflect the underlying shape of the object in a consistent manner.

In Figure 3.14 we provide some more detailed analysis of the specular reflection model. The left column of the figure shows the original grey-scale levels and the reconstructed Lambertian grey-scale levels in a transverse slice across the central specularity on the urn (shown in Figure 3.11). The different rows of Figure 3.15 show the results obtained with the Beckmann, Gaussian and Torrance-Sparrow models. The Gaussian model appears to give the best behavior below the specular intensity spike. In the right-hand column, we show the scatter plot of $\log P(\omega_{i,j} = S|\vec{M}_{i,j}^{(n)})$ versus $\cos\alpha$ for three different specular reflection models in the region around the central specularity. There are some significant differences here. In the case of the Gaussian model, the width of the specular component is wider than in the case of the remaining two distributions.

3.7.4 A Posteriori Probabilities

In Figures 3.15 to 3.19 we study the *a posteriori* probabilities for the specular reflectance using the three different distribution functions in the Bayesian framework. In each figure,

Figure 3.10: The porcelain urn.

Figure 3.11: Central highlight on the porcelain urn.

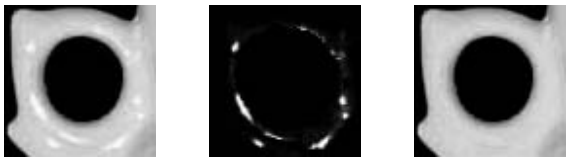

Figure 3.12: Specularities on the handle of the porcelain urn.

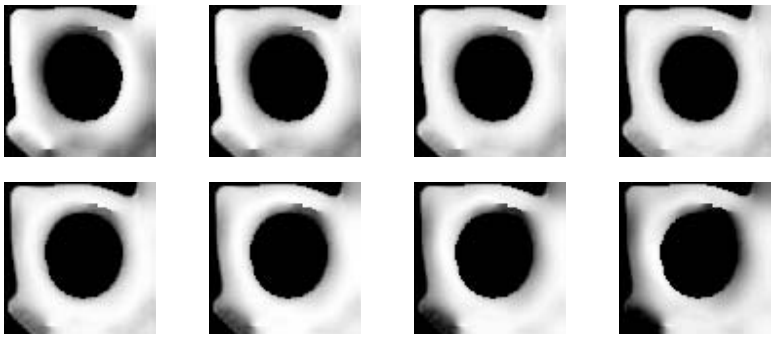

Figure 3.13: Re-illuminated images of the handle of the porcelain urn for incidence angles from -40° to $+40^\circ$.

the panels in the left-hand column are the results obtained using the Beckmann distribution function (Equation 3.20). The panels in the centre column are the results when the Gaussian distribution function (Equation 3.21) is used. Finally, the panels in the right-hand column show the results for the intensity-based distribution function using the Torrance-Sparrow model (Equation 3.22). In this experiment, the objects under study are the porcelain vase and the porcelain bear. The parameter of the Lambertian distribution function (Equation 3.23) is kept fixed at $\sigma_L = 0.1$, while the parameters of different specular distribution functions (σ_m, σ_α and σ_S) are varied.

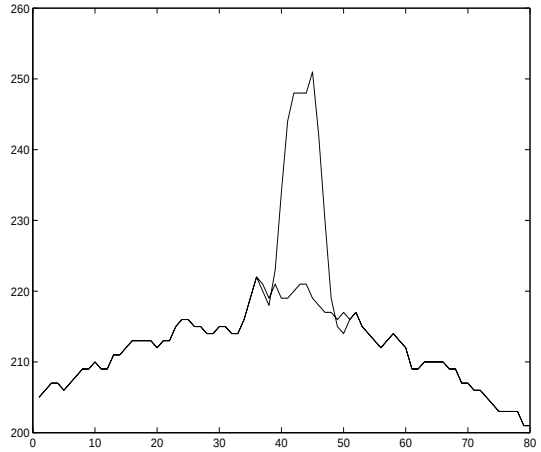

Raw and non-specular grey-levels (Beckmann)

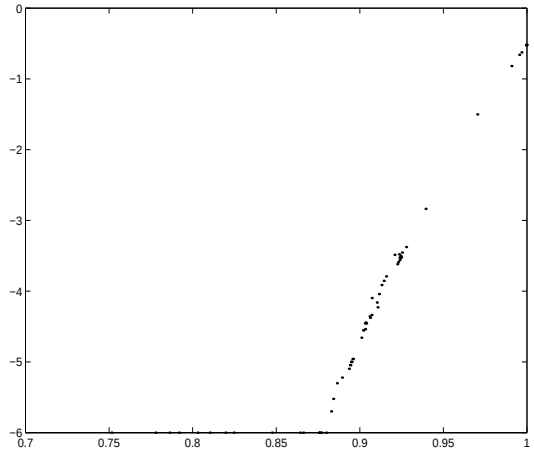

Scatter plot (Beckmann)

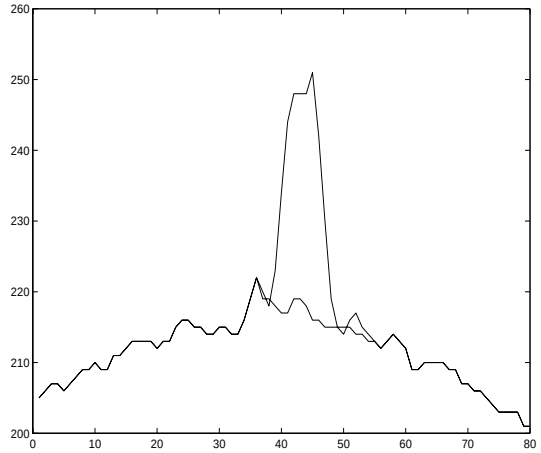

Raw and non-specular grey-levels (Gaussian)

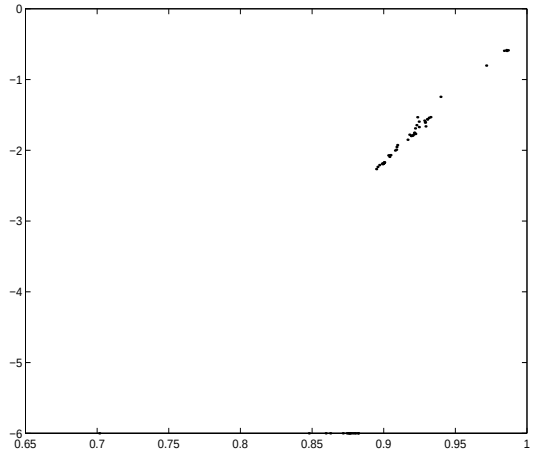

Scatter plot (Gaussian)

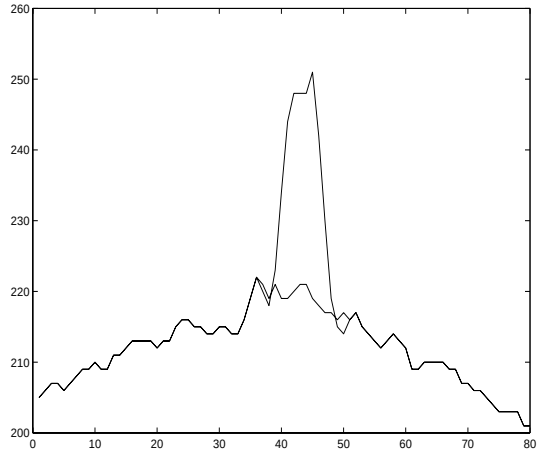

Raw and non-specular grey-levels (T-S)

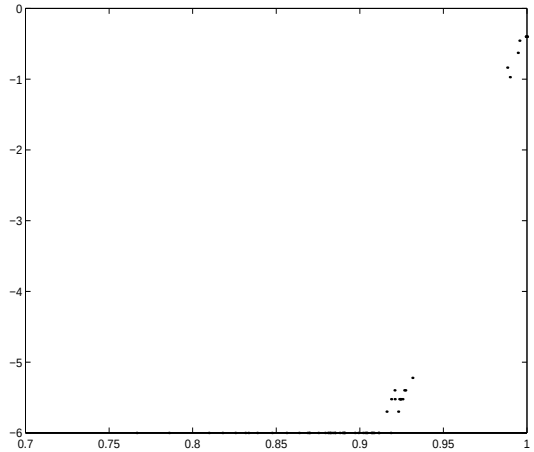

Scatter plot (T-S)

Figure 3.14: Results for the central detail on the urn (Figure 3.11) using different distribution functions: (a) the raw and matte grey-levels, (b) specular probability scatter plots.

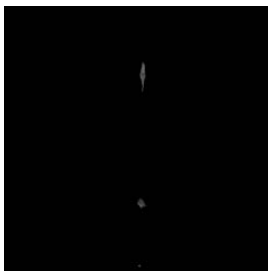

Spec. prob. (Beck., $\sigma_m = 0.1$)

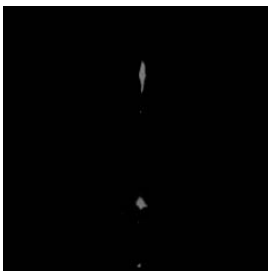

Spec. prob. (Gaus., $\sigma_\alpha = 0.1$)

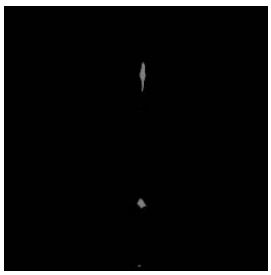

Spec. prob. (T-S, $\sigma_S = 0.1$)

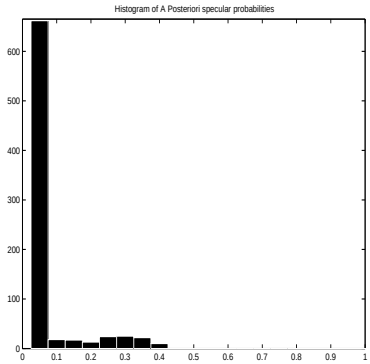

Beckmann ($\sigma_m = 0.10$)

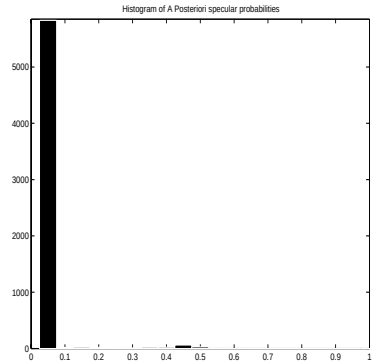

Gaussian ($\sigma_\alpha = 0.10$)

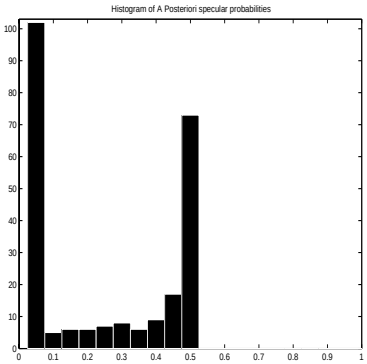

T-S ($\sigma_S = 0.10$)

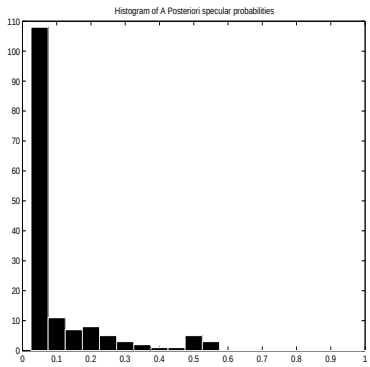

Beckmann ($\sigma_m = 0.05$)

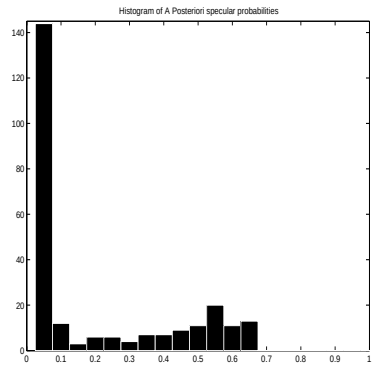

Gaussian ($\sigma_\alpha = 0.05$)

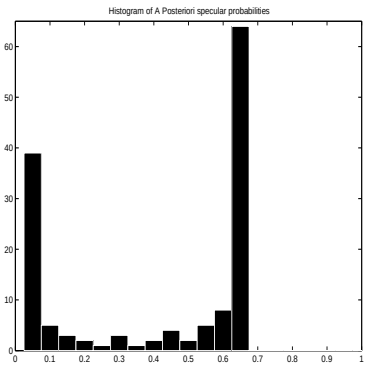

T-S ($\sigma_S = 0.05$)

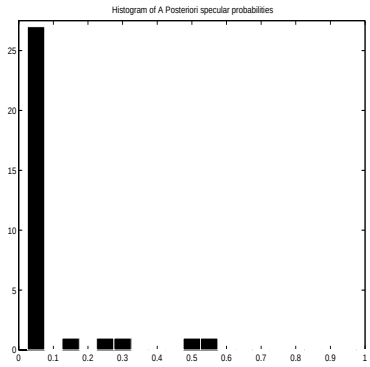

Beckmann ($\sigma_m = 0.02$)

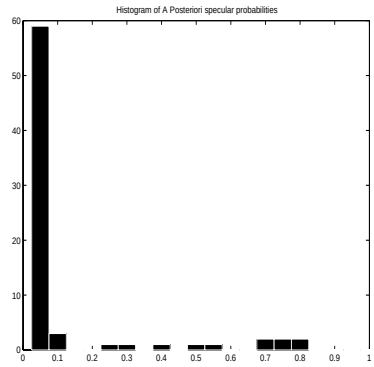

Gaussian ($\sigma_\alpha = 0.02$)

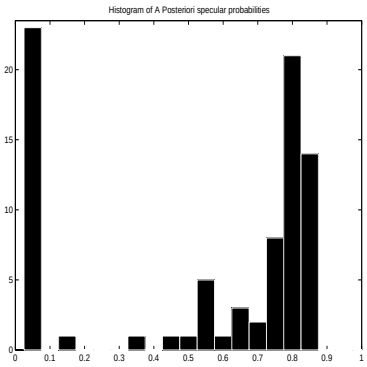

T-S ($\sigma_S = 0.02$)

Figure 3.15: Histograms of the (*a posteriori*) specular probabilities for the porcelain vase using different distribution functions and parameter values (σ_L is fixed at 0.1).

Spec. prob. (Beck., $\sigma_m = 0.1$)

Spec. prob. (Gaus., $\sigma_\alpha = 0.1$)

Spec. prob. (T-S, $\sigma_S = 0.1$)

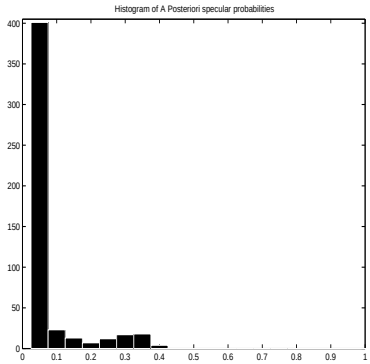

Beckmann ($\sigma_m = 0.10$)

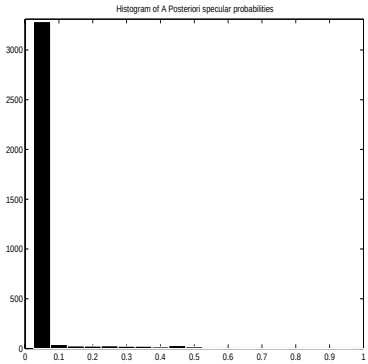

Gaussian ($\sigma_\alpha = 0.10$)

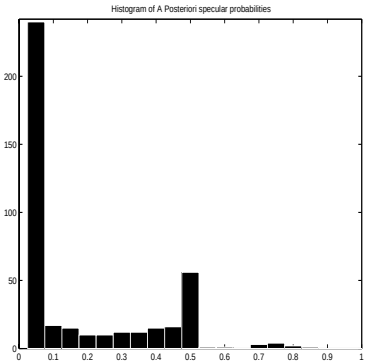

T-S ($\sigma_S = 0.10$)

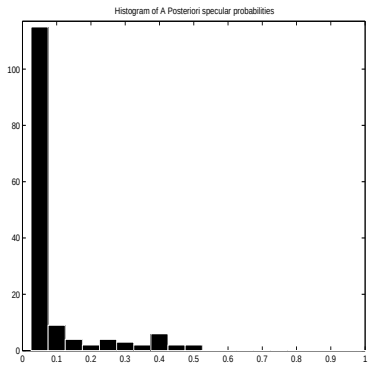

Beckmann ($\sigma_m = 0.05$)

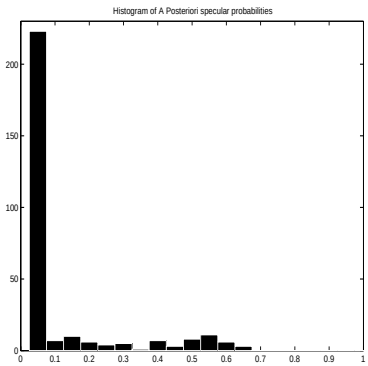

Gaussian ($\sigma_\alpha = 0.05$)

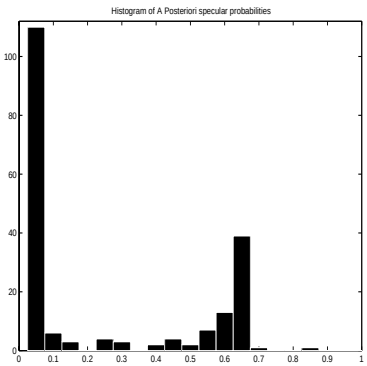

T-S ($\sigma_S = 0.05$)

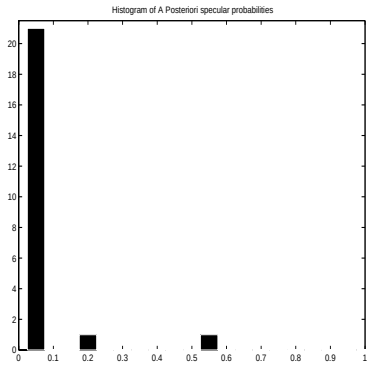

Beckmann ($\sigma_m = 0.02$)

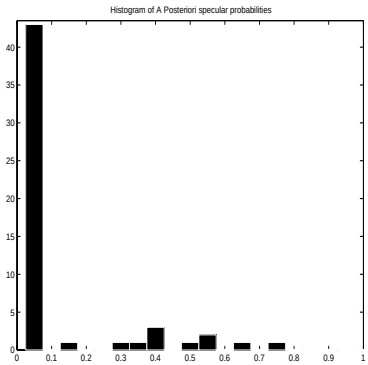

Gaussian ($\sigma_\alpha = 0.02$)

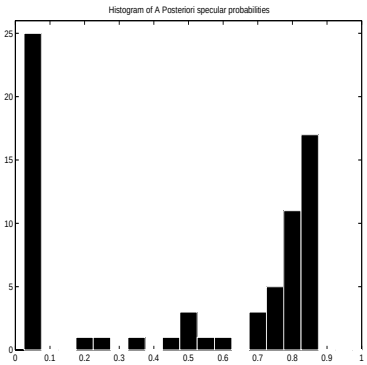

T-S ($\sigma_S = 0.02$)

Figure 3.16: Histograms of the (*a posteriori*) specular probabilities for the porcelain bear using different distribution functions and parameter values (σ_L is fixed at 0.1).

In Figures 3.15-3.16, from left to right, the panels in the top row show the specular probability images for $\sigma_m = 0.1$, $\sigma_\alpha = 0.1$ and $\sigma_S = 0.1$. Note that σ_m, σ_α and σ_S are the width parameters in the three distribution functions (Equations 3.20, 3.21 and 3.22). The second, third and fourth rows show the histograms of the non-zero specular probabilities for the three values 0.1, 0.05 and 0.02 used as the values of each distribution parameter. As we expect, the corresponding histograms for the vase (Figure 3.15) and the bear (Figure 3.16) are very similar in shape and behavior. As the parameter of the distribution decreases (from the second row to the bottom row), then so the section of high a posteriori specular probabilities increases. The Beckmann distribution is the least sensitive to the variations in the parameter, while the intensity-based distribution (T-S) is the most sensitive. Note that, by changing the value of σ_α or σ_S in the Healey-Binford distribution (Equation 3.21) or the Torrance-Sparrow intensity-based distribution (Equation 3.22), the width of the corresponding distribution changes while the Gaussian shape of the distribution remains unchanged. However, in the Beckmann distribution (Equation 3.20) changes in σ_m may result in changes in the shape of the distribution (Figure 3.3). The reason for this is that in addition to the exponential term which is present in the Gaussian distribution, the Beckmann distribution has the term $1/\cos^4\alpha$. It is this term which makes the control of the Beckmann distribution difficult.

In Figures 3.17, 3.18 and 3.19, we study the pixels assigned to the specular class and their correspondence with the histograms studied earlier. In each figure, we magnify a small patch of the corresponding image by an expansion factor of 10. In Figures 3.17 and 3.18, the specular highlights on the neck and the bulb of the porcelain vase are studied, while in Figure 3.19 the complex highlight on the left leg of the porcelain bear is studied. Here every distribution parameter, i.e. σ_m, σ_α, σ_S or σ_L, is set to 0.1.

In the top row of each figure, we show the original image of the magnified patch. The specular probability images are shown in the second row. The third row shows the reconstructed specular intensity using the Torrance-Sparrow model. Finally, the fourth row shows the non-specular intensity after performing specularity removal (Section 3.6). We use the grey-scale levels to show the probability values in the range [0, 1]. This means that

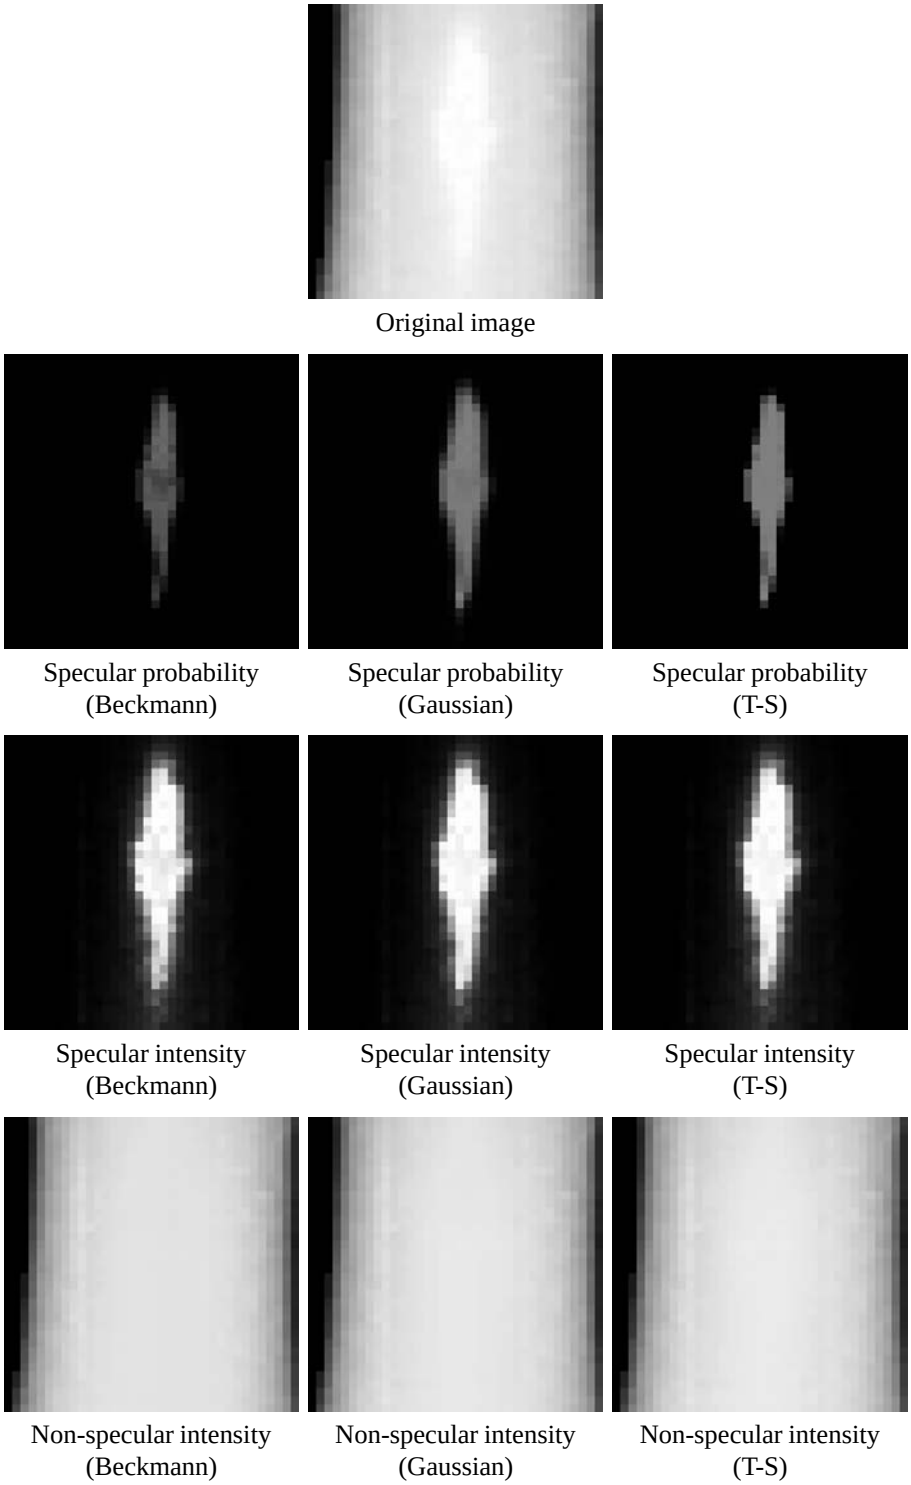

Figure 3.17: Results for highlights on the neck of the vase using different distribution functions (specular patches are enlarged to visualize grey-scale levels of the pixels).

the brighter a pixel on the specular probability image, the higher the specular probability.

Here, in addition to differences in the shape of the area labelled as the specular class, one

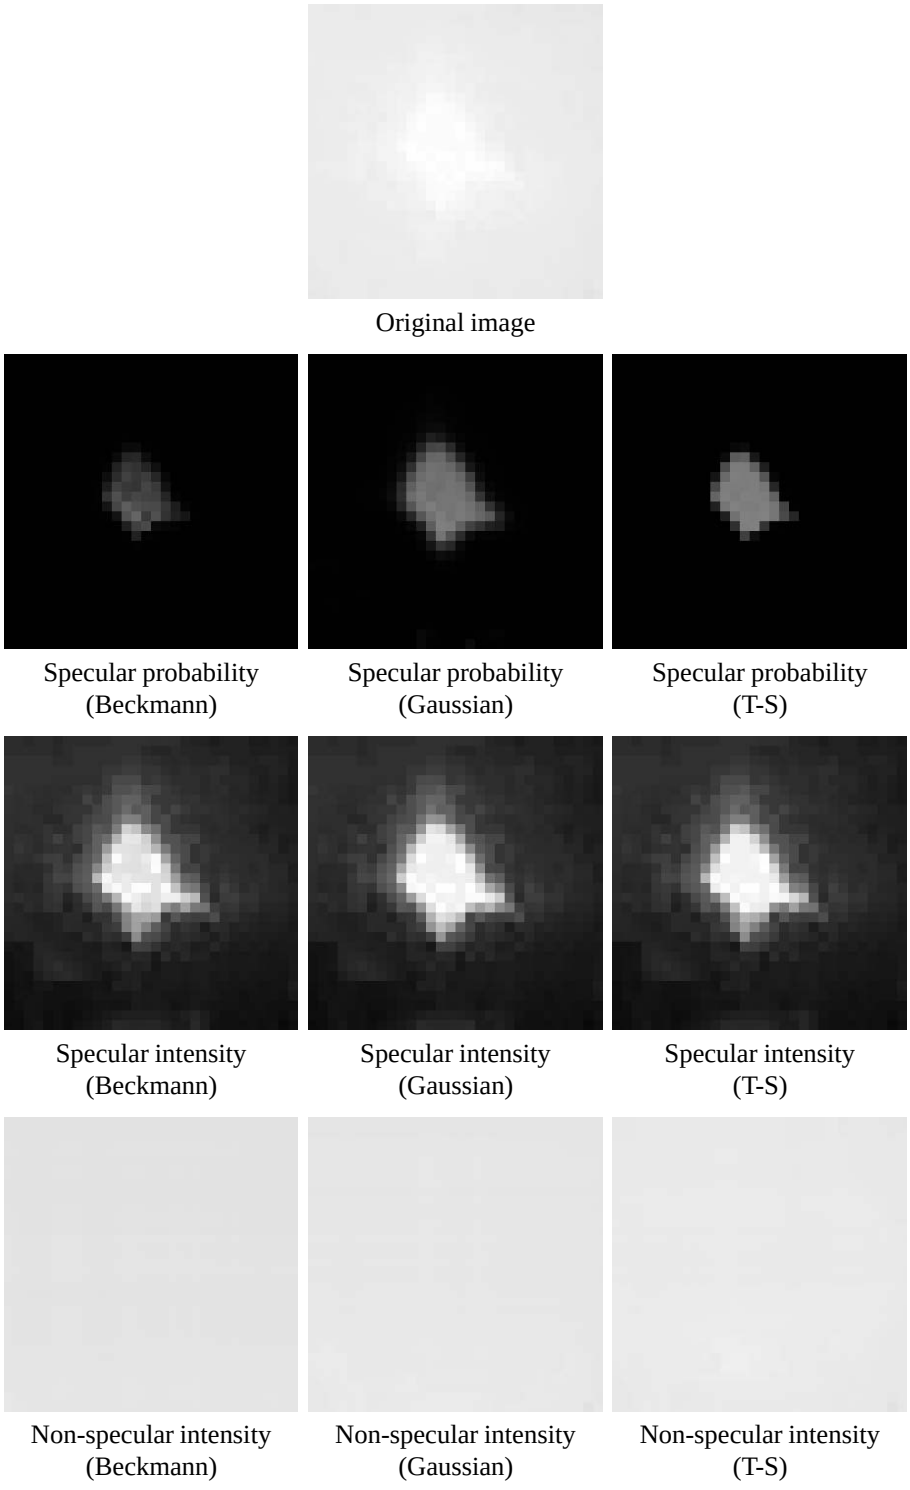

Figure 3.18: Results for highlights on the bulb of the vase using different distribution functions (specular patches are enlarged to visualize grey-scale levels of the pixels).

can also see differences in the brightness of the corresponding pixels in the three specular probability images of each sample. This is most marked for the specular patch of the

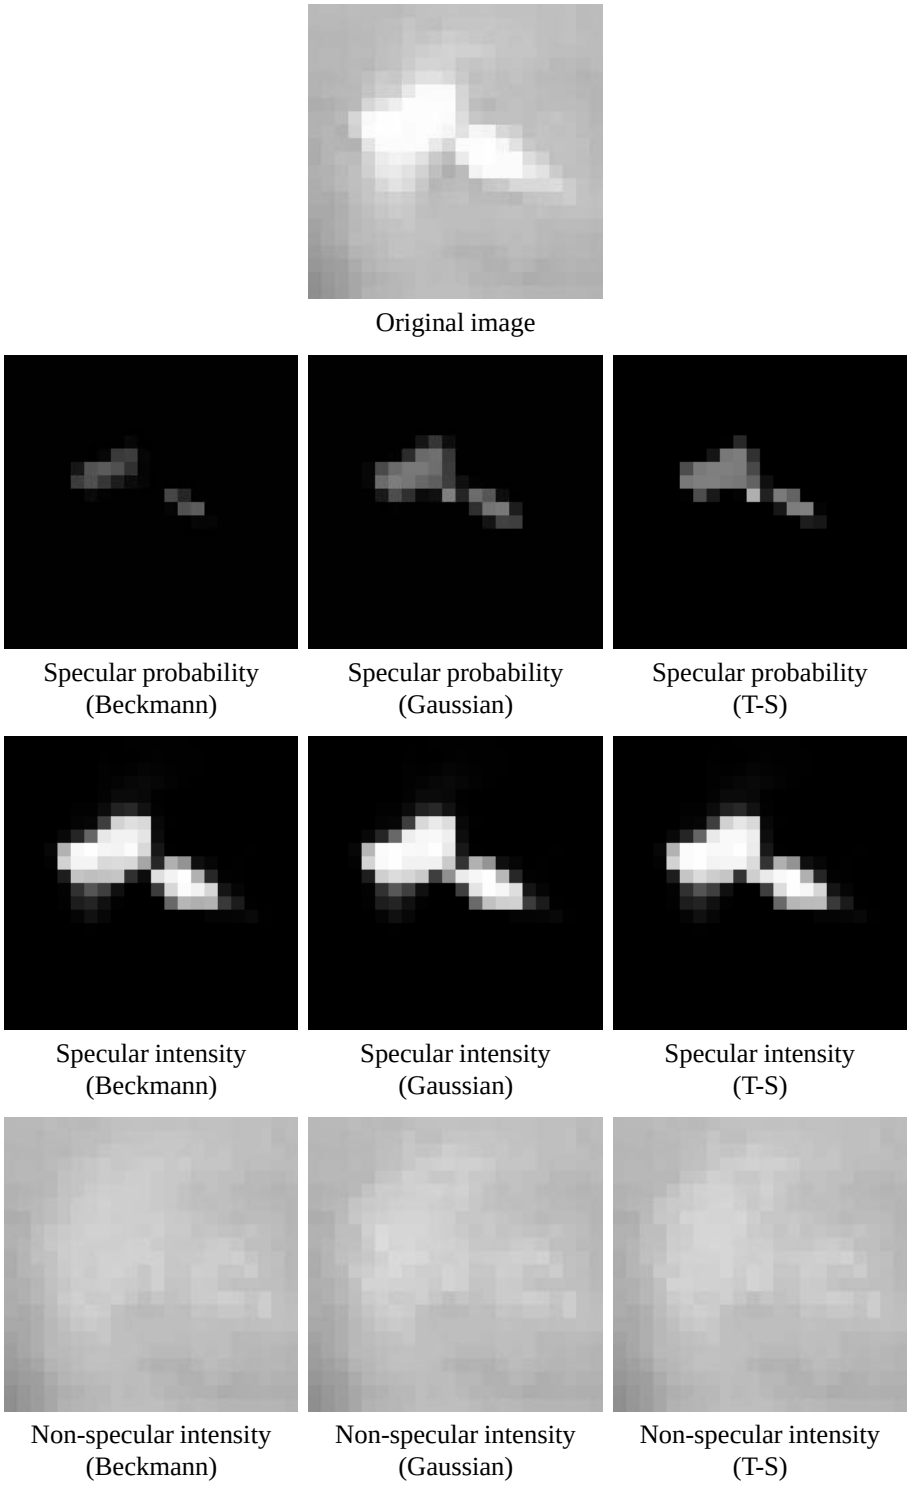

Figure 3.19: Results for highlights on the left leg of the bear using different distribution functions (specular patches are enlarged to visualize grey-scale levels of the pixels).

leg of the bear. Since the surface normals used in the reconstruction of specular intensity are posterior-mean surface normals (Equation 3.21), the balance between Lambertian

and specular *a posteriori* probabilities has an effect on the specular intensity too. Hence, although the same model (Torrance-Sparrow) is used to reconstruct the specular intensity images, they are slightly different when the different distribution functions are used in the Bayesian framework. Although the differences in the specular component could completely be transferred to the non-specular component during the highlight removal process, this has not happened and the differences between the non-specular components are smaller than differences between the specular components. Again, the differences can be observed more clearly in Figure 3.19 than in Figures 3.17 and 3.18.

3.7.5 Variable Light-Source Direction

We now illustrate the stability of the method as the light-source direction is varied. In the top row of Figure 3.20, we show four images of the porcelain vase as the light-source is moved. In the second row we show an enlarged detail of the bulb of the object which contains a clear specular reflection. In the third and fourth rows we show the matte and specular components extracted from these details. Again, the matte images are good and there appears to be no residual specular component. To underline this point, in Figure 3.21 we show the matrix of difference images for the raw data. Figure 3.22 repeats this matrix for the matte images obtained by specularity removal. The images in Figure 3.23 are the differences between Figure 3.21 and Figure 3.22. In the case of the raw images the main differences occur at the positions of the specular reflections and at the limbs of the object. The movement of the specularity across the surface is clear. In the case of the matte images, there are no residual artifacts in the proximity of the specularities and the only visible effects are at the limbs of the object. Finally, the only structures which persist in Figure 3.23 are the specularities.

To conclude the study of variable light-source direction, in Figure 3.24 we show the needle maps obtained from the matte images after specularity removal. In the top row we show the needle maps. The second row shows the difference fields for pairs of needle maps. The main structure is in the limb-regions which emerge from shadow as the light-

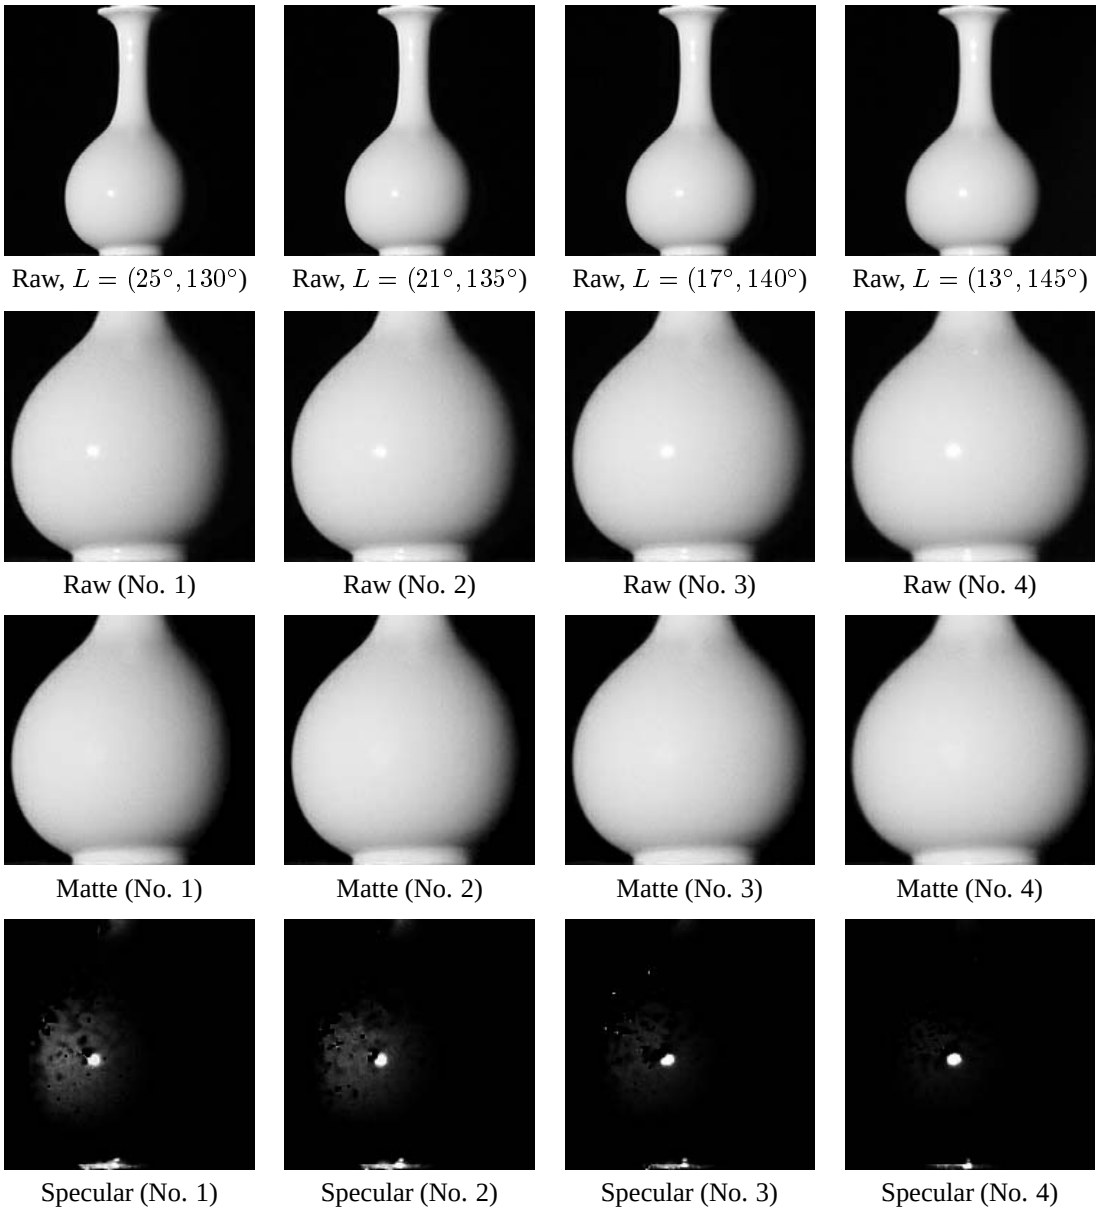

Figure 3.20: Removing specular highlights from images of the vase taken under different light-source directions, while the viewing direction is aligned with the z axis.

source moves. There are no significant differences in the proximity of the specularities.

3.7.6 Ground-Truth Data

We have also compared the surface normals with ground-truth data estimated from real-world imagery. Here we have studied a simple object, where the surface normal directions can be predicted from image plane measurements. To do this we have collected an image

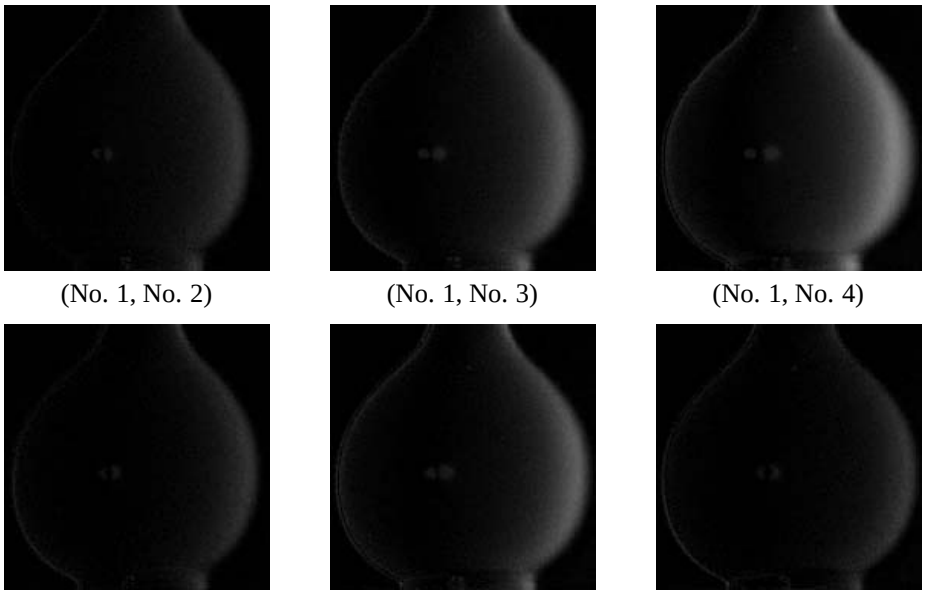

(No. 1, No. 2) (No. 1, No. 3) (No. 1, No. 4)

(No. 2, No. 3) (No. 2, No. 4) (No. 3, No. 4)

Figure 3.21: Pairwise difference images before specularity removal.

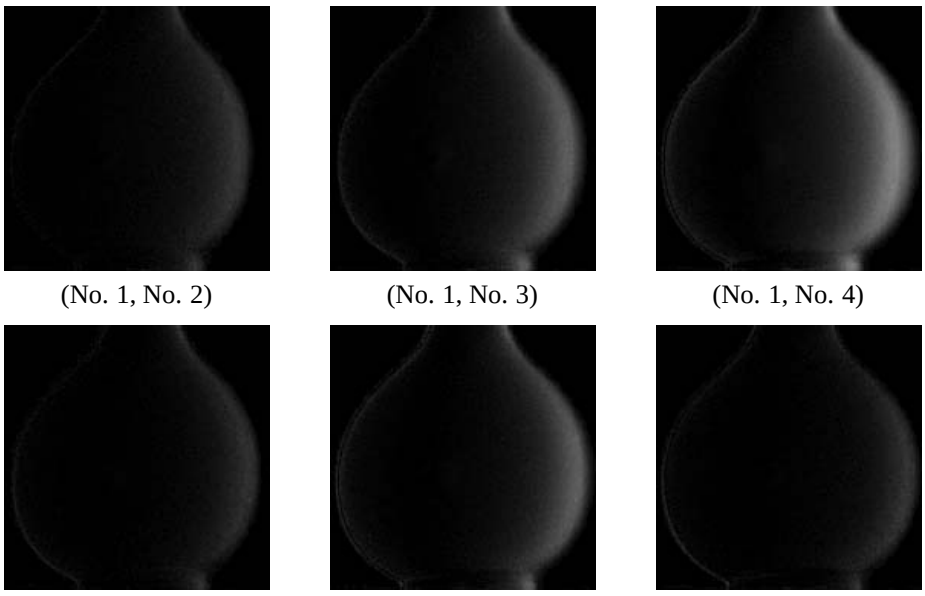

(No. 1, No. 2) (No. 1, No. 3) (No. 1, No. 4)

(No. 2, No. 3) (No. 2, No. 4) (No. 3, No. 4)

Figure 3.22: Pairwise difference images after specularity removal.

of a cylindrical object. We have aligned the cylinder so that its axis of symmetry is vertical (i.e. aligned with the y axis). If d is the measured diameter of the cylinder on the image, x_0 is the horizontal coordinate of the symmetry axis and x is the measured horizontal coordinate of the point on the cylinder, then the predicted unit surface normal at the measured point is $\vec{N}_{cyl}(x) = (\cos\gamma, 0, \sin\gamma)^T$ where $\cos\gamma = 2(x - x_o)/d$. From

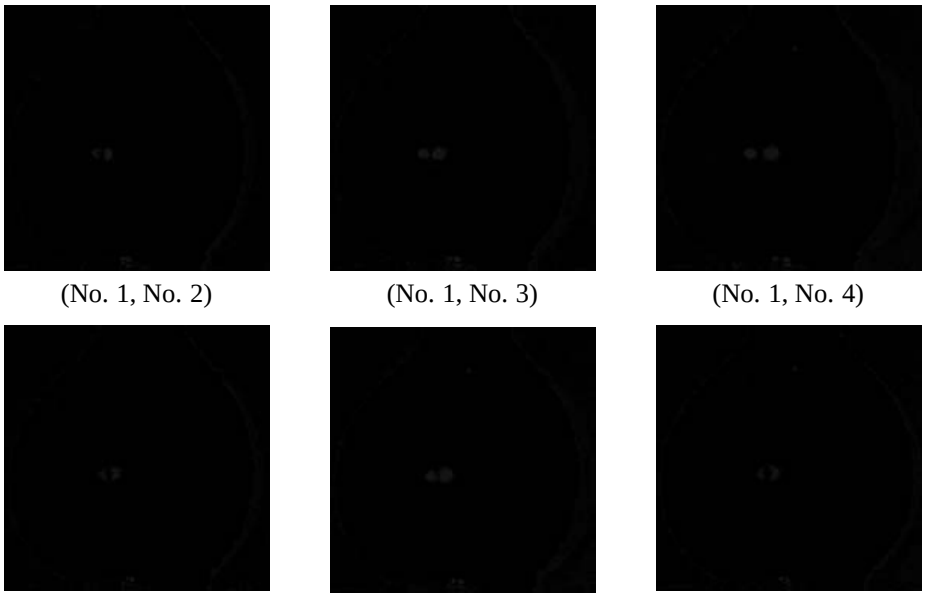

(No. 1, No. 2) (No. 1, No. 3) (No. 1, No. 4)

(No. 2, No. 3) (No. 2, No. 4) (No. 3, No. 4)

Figure 3.23: Residual differences between difference images of Figures 3.21 and 3.22.

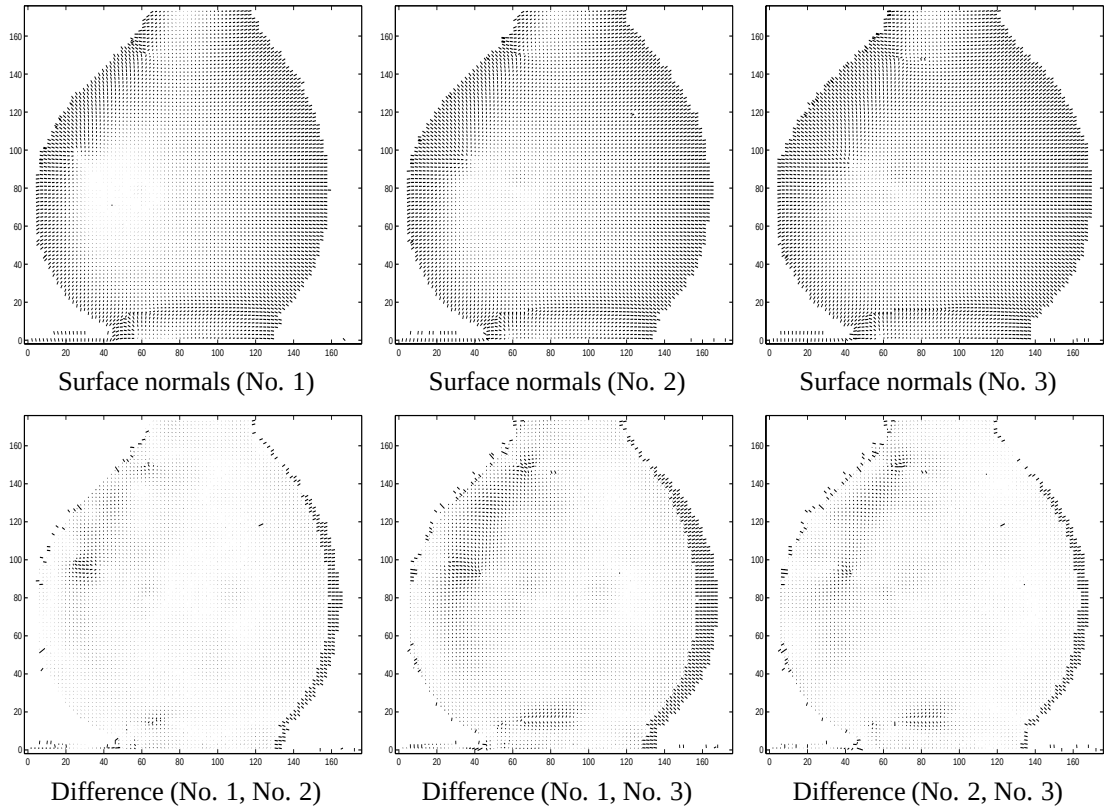

Surface normals (No. 1) Surface normals (No. 2) Surface normals (No. 3)

Difference (No. 1, No. 2) Difference (No. 1, No. 3) Difference (No. 2, No. 3)

Figure 3.24: Needle maps of (1) recovered surface normals for separated matte images of Figure 3.20; (2) corresponding fields of difference vectors.

the predicted surface normal, we can compute errors for the reconstructed surface normals at corresponding image locations.

In Figure 3.25 we show a study of imagery collected using a tea-cup. In the left-hand panel of the top row we show the original image. There is large specular reflection near the centre of the cup. The right-hand panel of the top row shows the predicted Lambertian intensity $E_{cyl}(x) = \vec{L}.\vec{N}_{cyl}(x)$. In the second row, the left-hand and right-hand images respectively show the Lambertian and specular reflectance components recovered using our new method. The Lambertian intensity component agrees well with the predicted value. In the third row the two panels show the needle maps extracted using the SFS method of Worthington and Hancock. The left-hand image is the result obtained when the method is applied to the Lambertian reflectance component, while the right-hand image is the result obtained when the method is applied to the raw image. In the bottom row of the figure we show difference fields between the two needle maps in the third row and the predicted surface normal directions computed using the simple geometrical model described earlier. In the case of the Lambertian reflectance component, there is no detectable structure in the difference field. However, when the raw image is used, then there is significant error in the region of the specular highlight.

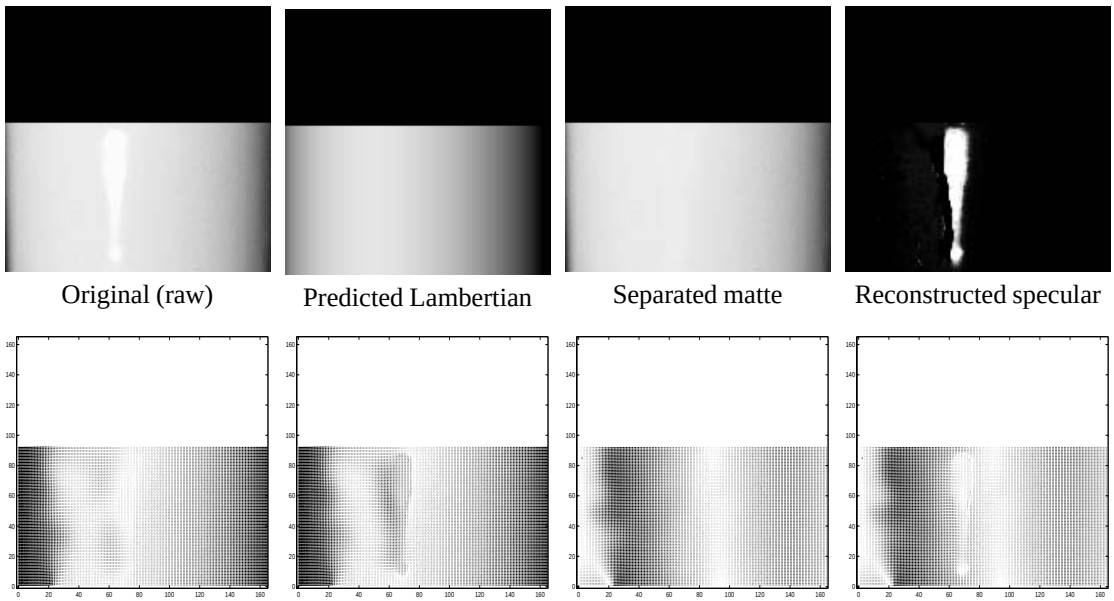

Figure 3.25: Surface normals for the raw and matte images compared to the ground-truth.

3.8 Conclusions

In this chapter, we have shown how to estimate the proportions of matte and specular reflectance from grey-level images of shiny curved objects. To do this we have developed a *maximum a posteriori* probability estimation framework for specular SFS. The algorithm uses separate probability density functions to model the angular distributions for the surface normals resulting from the specular and matte reflectance components. The method allows us to remove the specular reflectance component from shiny surfaces. This allows us to recover surface shape information which does not contain the curvature artifacts produced by local specular highlights. The method does not prove sensitive to the choice of specular reflection distribution model. The method has been demonstrated to be effective on a variety of real-world images.

There are a number of ways in which the work reported in this chapter can be extended and improved. In the first instance, the probabilistic framework for our study is provided by the iterated conditional modes algorithm. Although this proves convenient for our purposes, it has a number of well known shortcomings. These are primarily associated with local convergence. Hence, one direction for future research is to investigate the use of methods with known global convergence properties including mean-field annealing and simulated annealing. The second shortcoming of our work lies in the fact that the parameters of the reflectance distributions have been set by hand. However, Besag (Besag, 1986) suggests various neighborhood coding schemes that can be used for the purposes on online parameter estimation. Thirdly, we have used the probability distributions to model the physical variability of the surface normal directions, rather than their associated statistical uncertainties due to poor intensity measurements. Fourthly, and finally, our smoothness prior is extremely simple and leaves scope for considerable improvement.

Chapter 4

Non-Lambertian Diffuse Reflectance Models for Photometric Correction

In the previous chapter, we showed how we obtain diffuse component by removing specular component from the raw image of shiny surfaces. In doing this, we assumed that diffuse (matte or non-specular) component follows Lambert's law. However, it is well known that the reflectance properties of surfaces that are both rough and shiny depart significantly from the predictions of Lambert's law. These departures mean that standard Lambertian shape-from-shading (SFS) methods can not be applied directly to the analysis of rough and shiny surfaces. In order to overcome this difficulty, in this chapter, we consider how to reconstruct the Lambertian component when the object is illuminated in the viewing direction. To do this we make use of the diffuse reflectance models described by Oren and Nayar (Oren and Nayar, 1995), and by Wolff (Wolff, 1994a). However, and as it will be shown in the next chapter, this approach can be extended to other complex reflectance models.

Phenomenological models attempt to account for departures from Lambertian reflectance in a parametrically efficient manner. However, although these models provide a more accurate account of reflectance from rough and shiny surfaces, they do not prove tractable for estimating surface normal directions in SFS or other problems involving the analysis of surface characteristics or shape. Nonetheless, as we shall show in this chapter,

they do provide a route by which the Lambertian component can be recovered and used for the purposes of SFS. The approach used in our study is to construct a lookup table which allows the cosine of the incidence angle, and hence the Lambertian component, to be recovered from the measured brightness.

The work presented in this chapter is also motivated by the fact that several of the most successful existing methods for SFS are restricted in applicability to the Lambertian reflectance model and can not be adapted easily to work with the reflectance map. Hence, our correction process can be used as a simple tool to preprocess the input image data for such methods. Because the correction process delivers corrected information about the angle of incidence at every surface location, this allows the Lambertian SFS method to recover improved estimates of surface normal direction. In other words, our correction process allows existing Lambertian SFS methods to be applied to a greater variety of surfaces. It is also worth noting that Lambertian correction can be useful to assist in visualization, since the limb brightening due to roughness can result in loss of perceived surface relief.

There are a number of obstacles to the correction process. Firstly, the models depend on both the direction of the incident light and the viewer. Second, even when these two directions are identical, then the Wolff model and the Wolff-Nayar-Oren model are not tractable because of both the complexity of the Fresnel coefficient and the multiplicative nature of the correction. To overcome these problems, we construct a lookup table which allows the cosine of the incidence angle, and hence the Lambertian component, to be recovered from the measured brightness. It is interesting to note that in the case of the Oren-Nayar model, which involves an additive rather than multiplicative correction of Lambert's law, then the photometric correction can be performed analytically. To apply the reflectance models and perform Lambertian correction, we need to know the surface roughness and the index of refraction parameters appearing in the models. Here, we use approximate values of these parameters for similar surfaces obtained from tabulations in the literature. Although the fact that the light-source and viewing directions must be identical is a restriction, it does offer practical advantages. When the viewer and light-

source directions are identical, then we are able to illuminate most of the visible side of the object, and so, to recover the most complete shape from a single image of the object. Our experiments with synthetic data and also ground-truth data from real-world images confirm that the photometric correction process results in improved surface normal directions when a Lambertian SFS method is used.

The outline of this chapter is as follows. Section 4.1 discusses the three different reflectance models used in our analysis and explains how they may be simplified with a view to extracting corrected Lambertian component. In Section 4.2 we describe how the reflectance models outlined in Section 4.1 may be used to extract corrected Lambertian component. Section 4.3 details experiments on synthetic and real-world images. Finally, Section 4.4 offers conclusions and identifies avenues for future research.

4.1 Diffuse Reflectance Models

Lambert's law is the simplest model for diffuse reflectance from a surface. It has been widely used in computer vision to recover shape from shading patterns. If θ is the angle of incidence for the illuminant on the surface, the normalized radiance is $I_L(\theta) = \cos\theta$. There are a large number of situations in which there are significant departures from Lambert's law. This is the case close to the object occluding boundary for all illumination directions. For illuminant incidence angles larger than 50° relative to the viewer direction, there are also significant departures from Lambert's law both near the occluding boundary and over a large portion of the object bounded by the shadow boundary (Wolff et al., 1998). However, it should be stressed that the basis of Lambert's law is purely phenomenological, and the reflectance function does not attempt to model the physical processes of light scattering. The method does, however, provide a surprisingly good account of reflectance from certain surfaces (packed Magnesium Oxide for instance). Attempts have been made to justify the method using physics, and these include the use of the distribution of black-body radiation from glowing particles.

In this study we investigate whether these non-Lambertian reflectance models can

be used in conjunction with SFS as alternatives to Lambert's law for rough and glossy surfaces. Of course the classical way in which to deal with non-Lambertian surfaces is to follow Horn and construct a reflectance map as input to SFS (Horn, 1981). However, here we adopt a different approach, and explore whether we can use the reflectance models to extract corrected Lambertian component from the measured image brightness of non-Lambertian surfaces. This allows us to apply simple SFS algorithms, without the need to construct the reflectance map.

Unfortunately, since these non-Lambertian reflectance models depend on both the incidence and reflectance angles they prove to be intractable for arbitrary geometry. To overcome this problem, we restrict our attention to the case where the difference between light-source direction $\vec{L}$ and the viewing direction $\vec{V}$ is small, i.e. $\vec{L} \approx \vec{V}$. Although this is a restrictive assumption, it still results in methods which are of practical utility. Few would view SFS as a tool for analyzing objects in unstructured scenes and uncontrolled lighting. Instead, it may prove useful as a means by which to acquire surface shape under controlled lighting conditions. The method described in this paper is one that can be used to recover corrected Lambertian component under such conditions. For instance, Zhang and Shah (Zhang and Shah, 1997) have mentioned that when only one input image is used, in order to recover the shape as completely as possible, the image has to be taken with careful light-source placement to illuminate most of the object. Moreover, it is important to stress that the restriction does not weaken or approximate the contributions from the non-Lambertian component.

4.1.1 Oren-Nayar Model for Rough Surfaces

Oren and Nayar (Oren and Nayar, 1995) have developed a diffuse reflectance model for rough surfaces. They have used the roughness model proposed by Torrance and Sparrow (Torrance and Sparrow, 1967) that assumes the surface is composed of extended symmetric V-shaped cavities. Each cavity consists of two planar facets. The width of each facet is assumed to be small compared to its length. The roughness of the surface is specified

using a probability distribution function for the facet slopes. Further, it is assumed that the facet area is large compared to the wavelength of the incident light, and hence, geometrical optics can be used to derive the reflectance model. Finally, each facet is assumed to follow the Lambert's law. In its general form, the model takes into account complex geometrical effects such as shadowing, masking and interreflections between points on the surface. However, here we use their qualitative reflectance model in which interreflections are ignored to simplify matters. According to this model (Figure 4.1), for a point on a rough surface with a roughness parameter σ_a and illuminant incidence direction (θ_i, ϕ_i) and reflectance direction (θ_r, ϕ_r), the radiance is given by

$$L_r(\theta_i, \theta_r, \phi_r - \phi_i; \sigma_a) = \frac{\rho}{\pi} E_0 \cos\theta_i \{A + B \max[0, \cos(\phi_r - \phi_i)] \sin\alpha \tan\beta\} \quad (4.1)$$

where the parameters A and B, which are dimensionless, are only dependent on the surface roughness σ_a which measures the distribution of cavity wall slope angles, and is hence measured in degrees or radians (rd). The corresponding equations are:

$$A = 1.0 - \frac{0.5\sigma_a^2}{\sigma_a^2 + 0.33} \quad , \quad B = \frac{0.45\sigma_a^2}{\sigma_a^2 + 0.09} \quad (4.2)$$

The angles α and β are the maximum and minimum of the values of the zenith angles θ_i and θ_r, i.e. $\alpha = \max[\theta_i, \theta_r]$ and $\beta = \min[\theta_i, \theta_r]$. It is important to note that the model reduces to the Lambertian case when $\sigma_a = 0$. Finally, the parameter ρ represents the surface albedo (reflectivity) which is assumed to be constant.

Reflectance Geometries

In this section we consider the different geometries which give rise to distinct rough reflectance behaviors according to the Oren-Nayar model. When the light-source and the viewing directions are different, there are four possible cases which need to be considered. In Figure 4.2 we show the geometries for the different cases. In the figure we show two cones. The first of these has its axis in the light-source direction and has an apex

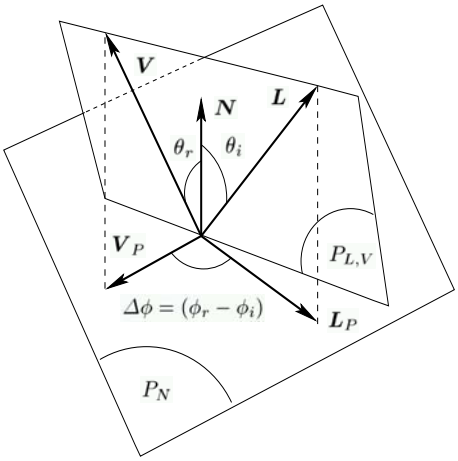

Figure 4.1: Typical geometry for the local tangent plane P_N: the normal $\vec{N}$ to the surface, the light-source direction $\vec{L}$, the viewing direction $\vec{V}$ and the corresponding angles.

angle of θ_i (the incidence cone). This is the cone of surface normal directions that satisfy Lambert's law (Worthington and Hancock, 1999). The second cone has its axis in the viewing direction and an apex angle of θ_r. We refer to the second cone as the emittance cone. The valid surface normal direction must fall on the intersection of the two cones.

In what follows we let $\theta = \theta_i$ and $\Delta\phi = \phi_r - \phi_i$. For simplicity, we also let $\Psi = \cos^{-1}(\vec{L} \cdot \vec{V})$ be the angle between the light-source and viewing directions. We also refer to $\max[0, \cos(\phi_r - \phi_i)]$ as the "Max" term. Furthermore, the vectors $\vec{L}_P$ and $\vec{V}_P$ are the projections of $\vec{L}$ and $\vec{V}$ onto the local tangent plane P_N (Figure 4.1).

The four different geometries are listed below:

Case 1: In Figure 4.2.a, $\theta_r = \theta_i + \Psi$ and the incidence cone is surrounded by the emittance cone. Here $\alpha = \theta_r$ and $\beta = \theta_i$, therefore the radiance equation is

$$I_D(\theta) = \cos\theta[A + B\sin(\theta + \Psi)\tan\theta] \tag{4.3}$$

Case 2: In Figure 4.2.b, $\theta_r = \theta_i - \Psi$ and the emittance cone is surrounded by the incidence cone. Here $\alpha = \theta_i$ and $\beta = \theta_r$, therefore

$$I_D(\theta) = \cos\theta[A + B\sin\theta\tan(\theta - \Psi)] \tag{4.4}$$

Note that in both cases 1 and 2, $\Delta\phi = 0$, so the "Max" term is replaced by 1.

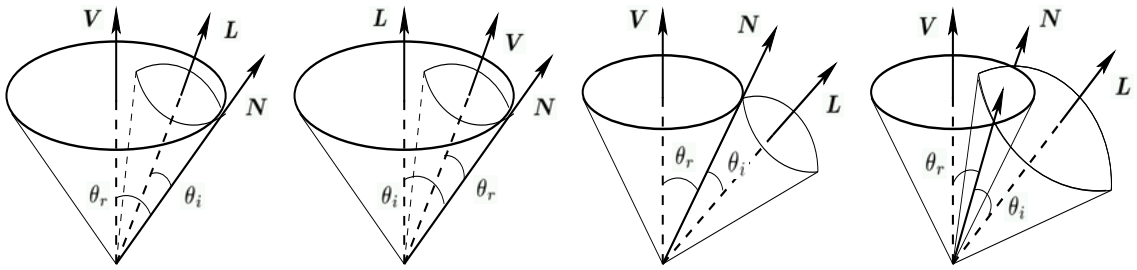

Figure 4.2: Reflectance geometries used to extract simplified equations from the Oren-Nayar model for the possible special cases; left to right: (a-d) Cases 1-4.

Case 3: In Figure 4.2.c, neither of the cones completely surrounds the other one and $\theta_r = \Psi - \theta_i$. Here $\Delta\phi = \pi$, and so the "Max" term is zero. Hence there is no need to identify α or β and we can write the radiance equation as

$$I_D(\theta) = A\cos\theta \tag{4.5}$$

Case 4: In Figure 4.2.d, $\theta_i + \theta_r > \Psi$ and the cones intersect each other. Here there are two possible surface normal directions, from which only one should be selected using additional constraints such as smoothness, curvature consistency or the proximity of occluding boundaries. For this case, we follow (van Ginneken et al., 1998) and write

$$\cos\Psi = \cos(\Delta\phi)\sin\theta_i\sin\theta_r + \cos\theta_i\cos\theta_r \tag{4.6}$$

Here, in general, it is not possible to find θ_r in terms of only θ_i. Similarly, α and β depend on the geometry. However, as a special case, when $\Psi = \pi/2$, $|\Delta\phi| < \pi/2$ and $\theta_i \geq \theta_r$, then the "Max" term is equal to $\cos(\Delta\phi)$, $\alpha = \theta_i$ and $\beta = \theta_r$. Since the left-hand side of Equation (4.6) is zero, i.e. $\cos(\Psi) = 0$, we can find θ_r in terms of θ_i and $\Delta\phi$:

$$\tan\theta_r = \frac{-\cos\theta_i}{\sin\theta_i\cos(\Delta\phi)} \tag{4.7}$$

Therefore, the radiance equation is given only in terms of incidence angle:

$$I_D(\theta) = A\cos\theta - B\cos^2\theta \tag{4.8}$$

For cases 1, 2 and 3 there is one possible surface normal $\vec{N}$ at the intersection line of the two cones. Moreover, the three vectors $\vec{L}$, $\vec{V}$ and $\vec{N}$ belong to the same plane $P_{L,V}$. These three cases cover any surface on which the curvature in one direction, i.e. the direction which is perpendicular to the plane $P_{L,V}$, is zero (e.g. parabolic or cylindrical surfaces). For case 4, on the other hand, there are two possible surface normal directions, and neither of these belongs to the same plane as $\vec{L}$ and $\vec{V}$. For instance, umbilic surfaces with indeterminate minimum and maximum curvature values, may fall into this category.

The Model for Photometric Correction

Here we use a simplified version of the Oren-Nayar model for the purposes of photometric correction, i.e. the process of recovering the Lambertian component. This method is applied when $\vec{L} \approx \vec{V}$. Under such conditions we can approximate the Oren-Nayar model by making the substitutions $\theta_r = \theta_i = \alpha = \beta = \theta$ and $\phi_i = \phi_r$. As a result, the simplified radiance (for $\rho = 1$) is

$$I_D(\theta) = A\cos\theta + B\sin^2\theta \tag{4.9}$$

Hence, the correction to Lambert's law is additive and proportional to $\sin^2\theta$. This term is greatest at the occluding boundary, and hence results in limb brightening. The constant B, and hence the roughness correction, vanishes when the roughness parameter $\sigma_a = 0$. Figure 4.3 shows the radiance versus the incidence angle for different values of the surface roughness parameter in Equation (4.9). The effect of increasing the roughness is to make the reflectance function flatter with incidence angle. Hence, the contrast between the limb and the remainder of the object is reduced.

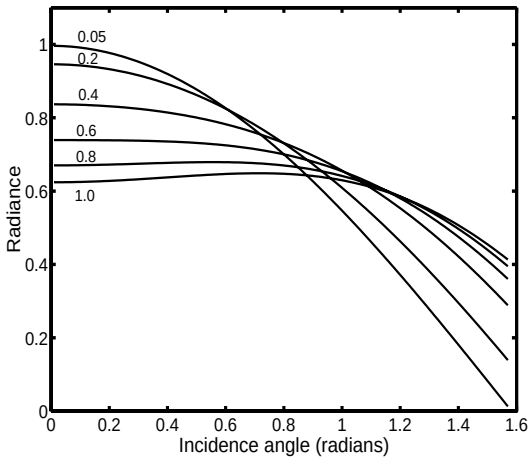

Figure 4.3: Surface radiance versus incidence angle using the simplified model of Oren-Nayar (Equation 4.9) for $\sigma_a = 0.05, 0.2, 0.4, 0.6, 0.8, 1.0$rd.

4.1.2 Wolff Model for Smooth Surfaces

Wolff has developed a physically motivated model for reflectance from smooth surfaces (Wolff, 1994a). The model accounts for subsurface refraction using a Fresnel attenuation factor, which modifies a Lambertian reflectance function in a multiplicative way. According to this model, the radiance is given by

$$L_r(\theta_i, \theta_r; n) = \varrho L_i \cos\theta_i [1 - f(\theta_i, n)]\{1 - f(\sin^{-1}[(\sin\theta_r)/n], 1/n)\} \qquad (4.10)$$

where n is the index of refraction of the dielectric medium. The attenuation factor is governed by the Fresnel coefficient $f(\alpha_i, r)$ where $0 \leq f(\alpha_i, r) \leq 1.0$. The angle α_i is the incidence angle. Note that α_i can be different from θ_i when light travels from a dielectric into air. The parameter r is the ratio of two index of refractions. Here the index of refraction for the air is equal to 1.0, while it is given by $n \neq 1.0$ for the dielectric. Specifically, when light is transmitted from air into a dielectric $r = n$ and $\alpha_i = \theta_i$. However, when transmission is from a dielectric into air, then $r = 1/n$ and $\alpha_i = \sin^{-1}[(\sin\theta_r)/n]$. The Fresnel coefficient is given by

$$f(\alpha_i, r) = \frac{1}{2}\frac{\sin^2(\alpha_i - \alpha_t)}{\sin^2(\alpha_i + \alpha_t)}\left[1 + \frac{\cos^2(\alpha_i + \alpha_t)}{\cos^2(\alpha_i - \alpha_t)}\right] \qquad (4.11)$$

where α_t is the transmission angle of light into the dielectric surface. This angle is given by Snell's law:

$$r = (\sin\alpha_i)/(\sin\alpha_t) \Rightarrow \quad \alpha_t = \sin^{-1}[(\sin\alpha_i)/r] \tag{4.12}$$

Using Fresnel attenuation and Snell refraction, the model accounts for how incident light and the distribution of subsurface scattered light behaves at a smooth air-dielectric surface boundary. The Wolff model deviates from the Lambertian form (i.e. $\cos\theta_i$) when the Fresnel reflection terms become significant. Almost all commonly found dielectric materials have an index of refraction, n, in the range of 1.4 to 2.0. For instance, for glass or silicon $n \approx 1.5$, for ceramic $n \approx 1.7$ and for titanium oxide or diamond $n \approx 2.0$. As a result the Fresnel function is weakly dependent upon the index of refraction for most dielectrics. Here, we do not consider the more complex case, which can be used to model paint, in which transparent particles of one dielectric are immersed in a carrier of another dielectric with a different index of refraction. The value of the scaling factor ϱ is very nearly constant over most incidence and reflectance angles (Wolff, 1994a). This scaling factor plays the same role as surface albedo in Lambert's law.

The Model for Photometric Correction

We can also develop a simplified version of the Wolff model when $\vec{L} \approx \vec{V}$. When $\theta_r = \theta_i = \theta$ the two Fresnel terms are identical and the radiance simplifies to

$$I_D(\theta) = \cos\theta[1 - f(\theta, n)]^2 \tag{4.13}$$

The Fresnel term has the effect of depressing the radiance for near-normal incidence. In Figure 4.4 we plot the radiance versus incidence angle using Equation (4.13) for different values of n. The greater the index of refraction, the greater the reduction in the near-normal reflectance.

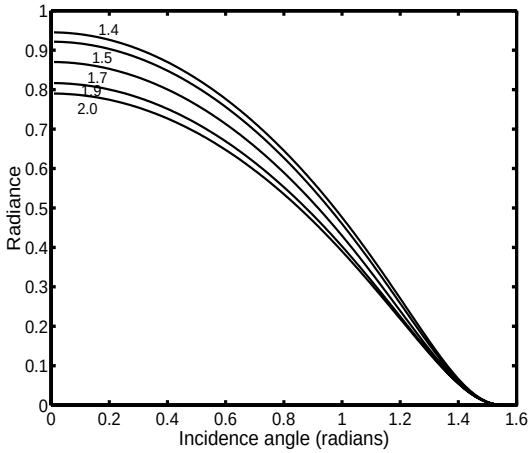

Figure 4.4: Surface radiance versus incidence angle using the simplified model of Wolff (Equation 4.13) for $n = 1.4, 1.5, 1.7, 1.9, 2.0$.

4.1.3 Combined Model for Intermediate Roughness

Unfortunately, there is a problem with the Oren-Nayar model for rough surfaces. As the roughness parameter decreases (i.e. $\sigma_a \rightarrow 0$), then the radiance is predicted to become Lambertian. However, this is a phenomenon that is supported neither by the Wolff model for smooth surfaces nor by empirical evidence (Wolff et al., 1998). Similarly, the Wolff model for smooth surfaces is restricted in applicability only up to a certain level of surface roughness. It is worth mentioning again the assumptions underpinning the two models. The Oren-Nayar model relies on the assumptions that the surface is composed of V-groove microfacets and that the reflectance from each microfacet is Lambertian. The Wolff model, on the other hand, assumes that there is subsurface optical scattering and air-dielectric boundary conditions. These effects are modelled by the Fresnel coefficient.

Based on this observation, Wolff et al. (Wolff et al., 1998) have commented on the fact that the two reflectance models, i.e. the Oren-Nayar model for rough surfaces (Oren and Nayar, 1995) and the Wolff model for smooth surfaces (Wolff, 1994a), are complementary in their applicability to surfaces with different roughness properties. Hence, they have combined the two models by making the assumption that each microfacet reflects according to the Fresnel model given by Equation (4.10). A simple approximation to the reflectance model that accounts for the observed empirical data for both rough and smooth

dielectrics is hence obtained by replacing A in Equation (4.1) with the quantity

$$C = A[1 - f(\theta_i, n)]\{1 - f(\sin^{-1}[\sin\theta_r/n], 1/n)\} \tag{4.14}$$

for an appropriate range of σ_a. This modification of the reflectance model is important since radiance from surfaces with intermediate roughness exhibits a combination of effects produced by both internal scattering and external roughness conditions.

The Model for Photometric Correction

We can also simplify the combined rough-smooth diffuse reflectance model to the case where $\vec{L} \approx \vec{V}$. The simplified radiance is

$$I_D(\theta) = A[1 - f(\theta, n)]^2 \cos\theta + B\sin^2\theta \tag{4.15}$$

Hence, the Fresnel term reduces the contribution from the Lambertian component at near-normal incidence. The $\sin^2\theta$ term is not affected, and results in limb brightening. In Figure 4.5 we show the behavior of this combined model. Here we show the radiance versus the incidence angle using Equation (4.15). The different curves are for different values of σ_a in the range $[0.01, 0.3]$ and the fixed value of $n = 1.7$. As the roughness increases, then so the proportion of radiance at large incidence angles increases, and the departure from the Lambertian case becomes more marked.

4.2 Photometric Correction Process

Our overall aim is to use the diffuse reflectance models discussed in the previous sections of this chapter to perform photometric correction and hence recover the Lambertian component. With corrected Lambertian images to hand, we can apply conventional SFS techniques to recover surface shape information. Here we use our recently developed SFS method outlined in Chapter 6. This method is an extension of the Worthington and Hancock algorithm (Worthington and Hancock, 1999). The surface normals are constrained

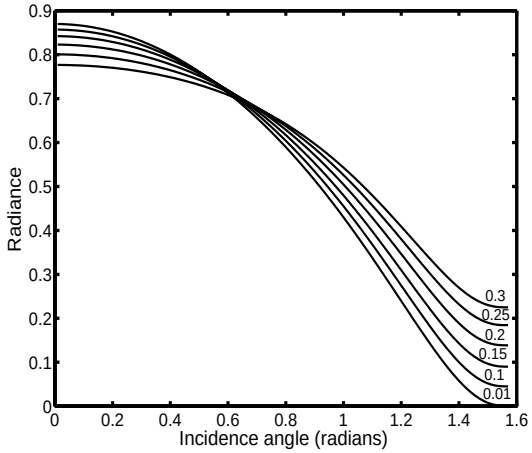

Figure 4.5: Surface radiance versus incidence angle using the simplified combined model (Equation 4.15) for $n = 1.7$ and $\sigma_a = 0.01, 0.1, 0.15, 0.2, 0.25, 0.3$rd.

to fall on a cone whose axis points in the light-source direction and whose apex angle is the inverse cosine of the corrected Lambertian component. However, rather than smoothing the field of surface normals using robust error kernels, we smooth the directions of the local Darboux frames using the field of principal curvature directions. Hence, we circumvent the problem of solving an under-constrained inverse reflectance equation to recover surface normal direction. It is this inverse problem that renders the non-Lambertian reflectance models intractable. By utilizing the field of principal curvature directions we recover more accurate surface normals both at parabolic locations on the surface and at the locations of surface discontinuities (e.g. intersections between shapes).

4.2.1 Analytic Solution

The Oren-Nayar model is amenable to an analytic solution under the following conditions. Suppose that at each image pixel (i, j) the diffuse (non-specular) radiance $I_D(i, j)$ consists of two components. The first of these is a Lambertian component $A \cos\theta$. The second is the non-Lambertian component $B \sin^2\theta$ which takes on its maximum value where $\theta = \pi/2$, i.e. close to the occluding boundary. To perform photometric correction, we proceed as follows. At every pixel location, we use Equation (4.9) to estimate $\cos\theta$

using the measured radiance and solving the following quadratic equation for $\cos\theta$.

$$-B\cos^2\theta + A\cos\theta + B - I_D(i,j) = 0 \tag{4.16}$$

The solution is

$$\cos\theta = \frac{A \pm \sqrt{A^2 - 4B[I_D(i,j) - B]}}{2B} \tag{4.17}$$

Since the Lambertian component is proportional to $\cos\theta$, this provides a means by which we can perform direct photometric correction and remove the effects of rough surface reflectance. The problems with using Equation (4.17) to estimate the Lambertian component are twofold. First, there are two possible solutions. Second, the method is singular when the argument of the square root is negative or when the constant B vanishes. The first problem can be overcome by taking the sign which results in a value of $A\cos\theta$ which is closest to the measured radiance I_D (in the majority of cases this involves taking the solution associated with the minus sign).

4.2.2 Lookup Table Solution

Although the Oren-Nayar model is analytically tractable for the purposes of Lambertian correction, both the Wolff model and the combined model of Wolff-Nayar-Oren are intractable to analytic solution (i.e. the problem of obtaining $\cos\theta$ in terms of I_D). For all three models, we therefore adopt a lookup table approach as a practical alternative. To do this we tabulate $\cos\theta$ as a function of the computed radiance from each model. Since the Lambertian component is proportional to $\cos\theta$, this allows us to perform Lambertian correction using the measured radiance values. In practice, the larger the number of incidence angles tabulated ($0 \leq \theta \leq \pi/2$), the more precise the correction process.

For the Oren-Nayar model, using an approximate value of the surface roughness $\sigma_a = \sigma_0$, the lookup table is computed from Equation (4.9). For the Wolff model, the radiance values are almost identical for most dielectrics. This applies when the index of refraction n is in the range $[1.4, 2.0]$. To give some idea of the physical significance of

this restriction, glass has a refractive index of about 1.5 while diamond has a refractive index which is even larger than 2. Hence, we can use the centre-range value (i.e. $n = 1.7$) when we compute the lookup table using Equation (4.13). Similarly, for the combined model, the lookup table can be computed using the approximate values of σ_a and n from Equation (4.15).

By assuming that $\vec{L} \approx \vec{V}$, all three models can be made amenable to our lookup table approach. The method is usable in practice since, like Lambert's law, the brightness decreases monotonically with increasing incidence angle (as shown in Figures 4.3-4.5). As a result the reflectance functions appearing in Equations (4.9), (4.13) and (4.15) are injective and invertible. Each measured brightness value I_0 is related to a single value of incidence angle θ_0 and hence to a single corrected Lambertian component $\cos\theta_0$.

4.3 Experiments

In this section we provide some experimental results for both synthetic and real-world objects when the object under study is illuminated in the viewing direction. The real-world images used in our experiments have been captured using an Olympus 10E camera. Each surface has been imaged under controlled lighting conditions in a darkroom. The objects have been illuminated using a single collimated tungsten light-source. For shiny surfaces (e.g. porcelain objects) both specular and diffuse components are present. The specular component must therefore be removed before we perform our photometric correction on the diffuse component. To remove the specular component, we use the method that we have previously described in Chapter 3. However, for matte surfaces (e.g. terra-cotta objects) the diffuse component is the dominant one. Hence, specular subtraction is not necessary. To recover surface normal directions, we use the Lambertian SFS method which is outlined later in Chapter 6.

4.3.1 Synthetic Data

In this section, we explore the use of the Lambertian correction process on synthetic images with known ground-truth surface normal directions. The surface studied here consists of a large elliptical dome surrounded by four smaller ones. In Figure 4.6.a we show the elevation or height data for the synthetic surface, while Figure 4.6.b shows the field of ground-truth surface normals. In Figure 4.7, we show the result of rendering the synthetic surface using the reflectance models studied in this chapter. From left to right, the panels show the images obtained using Lambert's law, the Wolff model ($n = 1.7$), the Oren-Nayar model ($\sigma_a = 20^\circ$), and the combined model ($n = 1.7$ and $\sigma_a = 20^\circ$).

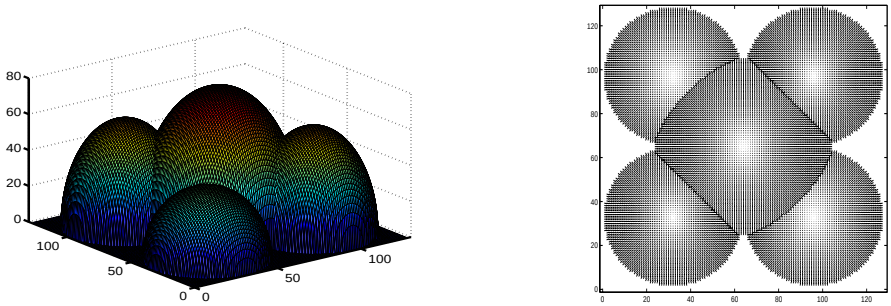

Figure 4.6: Synthetic surface: (a) elevation data (b) field of surface normal directions.

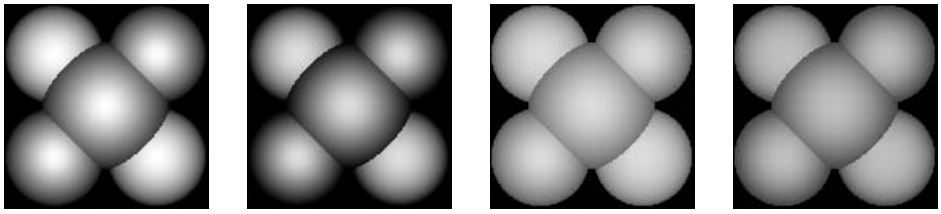

Figure 4.7: Images of the synthetic surface of Figure 4.6 rendered using the models of (a) Lambert, (b) Wolff ($n = 1.7$), (c) Oren-Nayar ($\sigma_a = 20^\circ$), and (d) Wolff-Nayar-Oren ($n = 1.7$ and $\sigma_a = 20^\circ$).

Applying Lambertian SFS to Non-Lambertian Images

To illustrate the errors introduced by using the incorrect reflectance model, we apply our Lambertian SFS method, which will be presented in Chapter 6, to each of the images of Figure 4.7. In the top row of Figure 4.8, we show the recovered fields of surface normals. In the bottom row of the figure, we show the difference between the estimated surface

normals and the ground-truth surface normals. The results are summarized in Figure 4.9, where we show scatter plots of the estimated zenith (top row) and azimuth (bottom row) angles of the estimated surface normals versus their ground-truth. From the zenith angle plots, it is clear that the greatest error occurs when the Oren-Nayar and Wolff-Nayar-Oren models are used. In the case of the Wolff model, the error is only significant when the zenith angle exceeds $60°$. As expected, there is zero error in the Lambertian case, since the geometric SFS method imposes agreement between the zenith angles as a hard constraint. In the case of the azimuth angle plots the errors are due to the effects of the surface normal smoothing process. Some significant errors occur when the Wolff model is used. The errors resulting from the use of the Lambertian, Oren-Nayar and combined models are comparable.

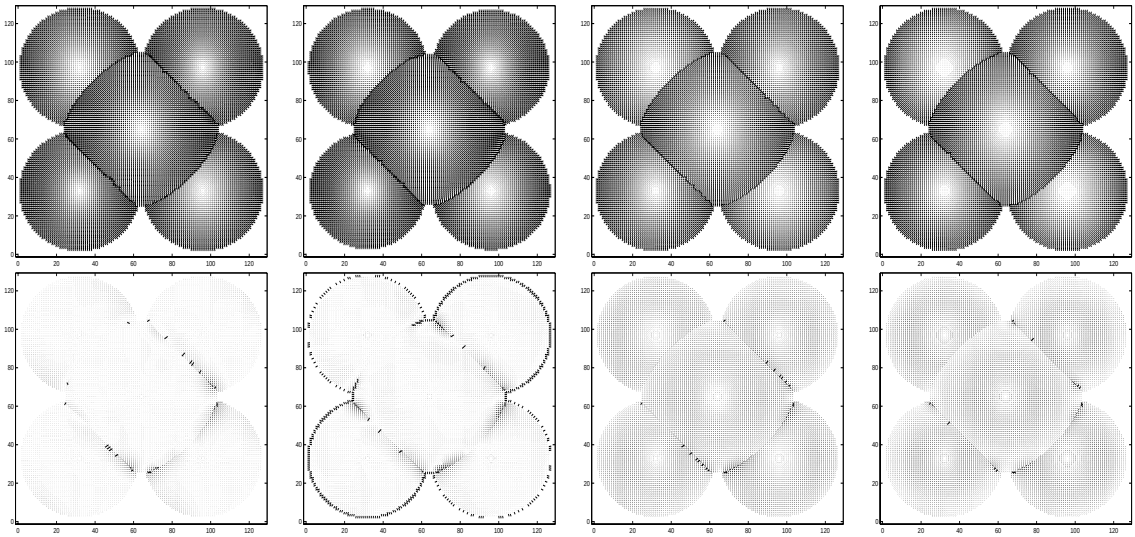

Figure 4.8: Applying Lambertian SFS (Chapter 6) to images of Figure 4.7: (1) recovered surface normals; (2) difference between the recovered and ground-truth surface normals.

To investigate the effect of the smoothing process in more detail, in Figure 4.10 we show the total-mean-square (TMS) error between the estimated and ground-truth surface normals as a function of iteration number. If $\vec{N}_{GT}(i,j)$ is the field of unit ground-truth surface normals and $\vec{N}_E(i,j)$ is the field of estimated unit surface normals, then the TMS error is defined to be $TMS = \Sigma_{i,j}\left[\| \vec{N}_{GT}(i,j) - \vec{N}_E(i,j) \|\right]$. As expected, the Lambertian model gives the smallest error. Of the three non-Lambertian models it is the Wolff

model that gives the lowest error and the Oren-Nayar model that gives the largest error. The combined model gives intermediate error. This experiment shows that when a Lambertian SFS method is used, the input image should be a Lambertian image rather than an image following other reflectance models. However, it also shows that it is important to find the appropriate correction model according to the surface roughness properties.

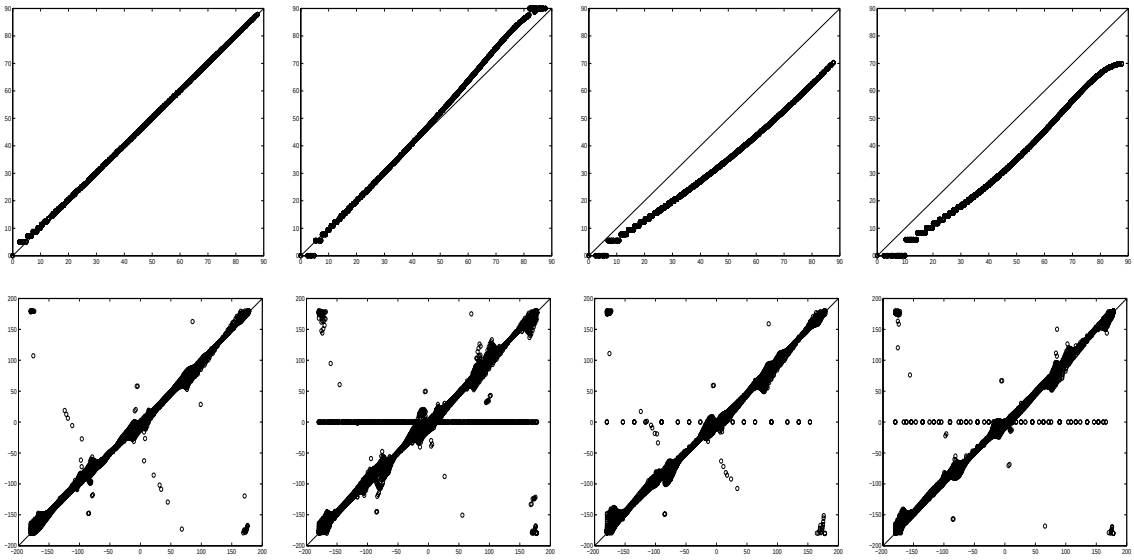

Figure 4.9: (1) Zenith angles and (2) azimuth angles of the recovered surface normals versus the corresponding ground-truth angles.

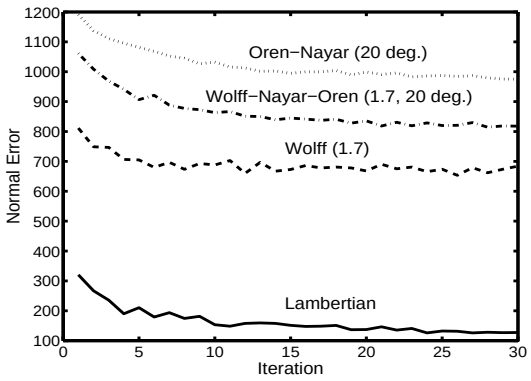

Figure 4.10: TMS surface normal error versus iteration number obtained by applying Lambertian SFS (Chapter 6) to the images of Figure 4.7 generated using different models.

Recovering Corrected Lambertian Images

In Figure 4.11, we illustrate the effect of applying the correction process to the synthetic images. In the top row, the panels from left to right show the corrected Lambertian images obtained by applying the table lookup method to the synthetic images generated using the Wolff model, the Oren-Nayar model and the combined model. In the bottom row, the panels show the difference between the corrected images and the original images (the darker the image location, the greater the difference). In all three cases, the effects of the correction process are greatest near the occluding limb of the object. The differences between the corrected images and the ground-truth Lambertian image are decreased. The values of the parameters used for constructing the lookup tables are different from those used when creating the images of Figure 4.7. Our aim in doing this is to illustrate that even using imprecise values of the surface parameters, the correction process results in improved surface shape recovered using the SFS method. The lookup table parameters are $n = 2.0$ in the Wolff model, $\sigma_a = 10^\circ$ in the Oren-Nayar model, and, $n = 1.7$ and $\sigma_a = 10^\circ$ in the combined model.

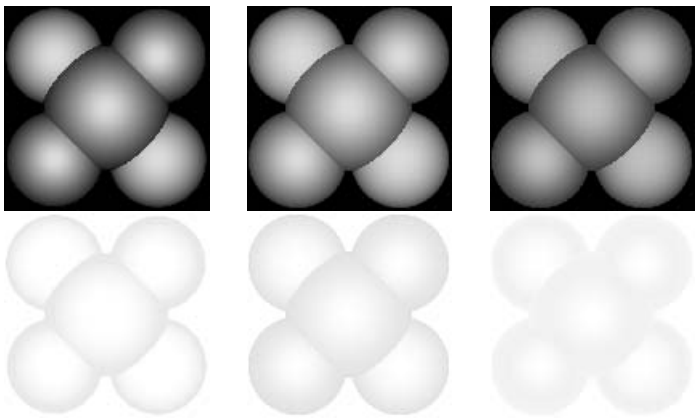

Figure 4.11: (1) The corrected images obtained by applying the table lookup method to the images of Figure 4.7 using the (a) Wolff, (b) Oren-Nayar, and (c) combined models; (2) difference images.

In Figure 4.12, we repeat this analysis by adding noise to the non-Lambertian images of Figure 4.7. We have added noise to 10% of the image pixels. The noise is Gaussian with zero mean and standard deviation equal to 10% of the original image brightness.

These noisy images are shown in the top row of the figure. The correction process is applied with the same parameter values as used in Figure 4.11. The corrected images are shown in the second row of Figure 4.12. Finally, the corresponding difference images are shown in the bottom row of the figure. In each case the differences are greatest at the most steeply inclined locations on the surface.

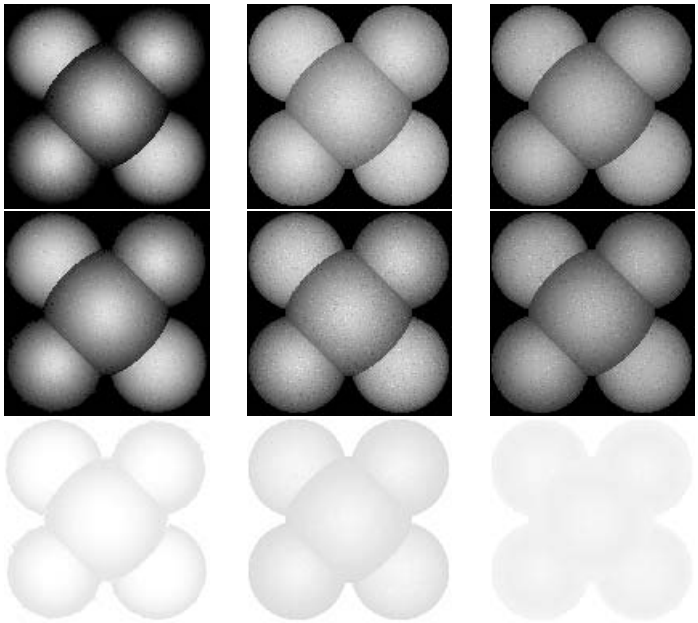

Figure 4.12: (1) Images obtained by adding Gaussian noise to images of Figure 4.7; (2) corrected images obtained by applying the table lookup method to the noisy images using the models of (a) Wolff, (b) Oren-Nayar, and (c) Wolff-Nayar-Oren; (3) difference images.

Applying Lambertian SFS to Corrected Lambertian Images

In Figure 4.13 we illustrate the effect of the correction process on the arc-cosine (ACS) error for the estimated surface normals. Each panel in the plot shows the image of $ACS = \cos^{-1}(\vec{N}_{GT}.\vec{N}_E)$, where $\vec{N}_{GT}$ is the unit ground-truth surface normal and $\vec{N}_E$ is the estimated unit surface normal at the corresponding image location. The panels in the top row show the ACS error images when the SFS algorithm is applied to the uncorrected images generated with the Wolff, Oren-Nayar and combined models. In the second row, we show the ACS error images obtained when the SFS algorithm is applied to the corrected images. The third and fourth rows of the figure repeat this analysis with the

noisy images. The main effects to note from the panels are as follows. First, the greatest error occurs at the cusp between the surfaces, where the SFS algorithm results in some oversmoothing of the field of surface normals. Second, while the errors are greatest at the object boundary when correction is not applied, the distribution in error across the surface is relatively uniform after correction. Finally, the effect of noise, is to increase the errors near the umbilic points at the crest of the domes, and to increase the error near the cusp-lines.

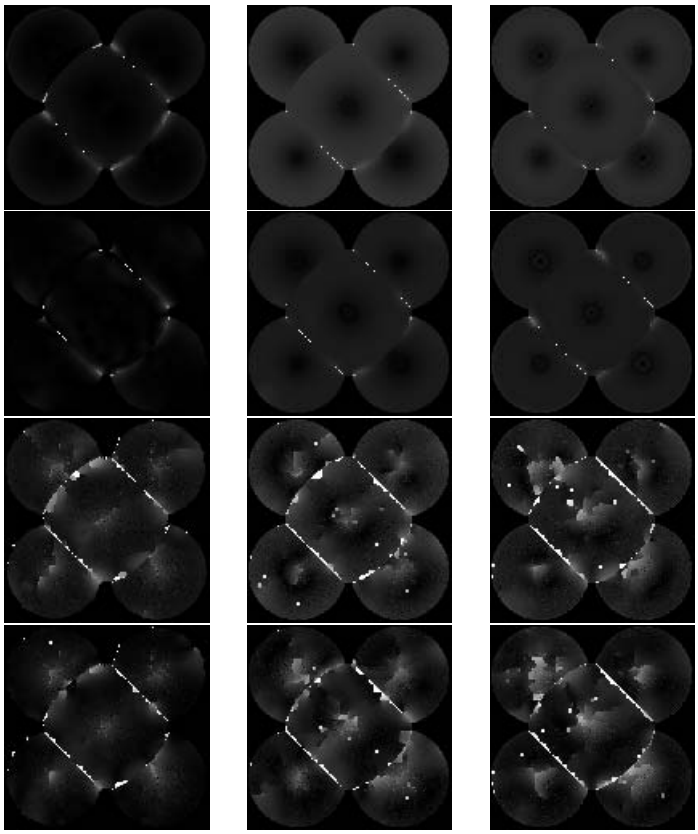

Figure 4.13: Images of arc-cosine (ACS) error between the ground-truth and the estimated surface normals corresponding to the (1) uncorrected images of Figure 4.7, (2) corrected images of Figure 4.11, (3) uncorrected and (4) corrected noisy images of Figure 4.12.

To analyze the behavior of the algorithm in more detail, in Figure 4.14 we show plots of the TMS error between the estimated and ground-truth surface normals as a function of iteration number. The plots in the top row of the figure are for the noise free images, while those in the bottom row are for the noise corrupted images. In each plot the dashed curve is the result obtained when the SFS method is applied to the uncorrected image,

while the solid curve is the result obtained using the corrected image. In each case the error is much lower when the correction process is applied. It is important to note that in the top row we use 30 iterations, while in the bottom row we use an extra 90 iterations to allow SFS to smooth away the effects of noise as much as possible. However, the results are encouraging even without these extra iterations.

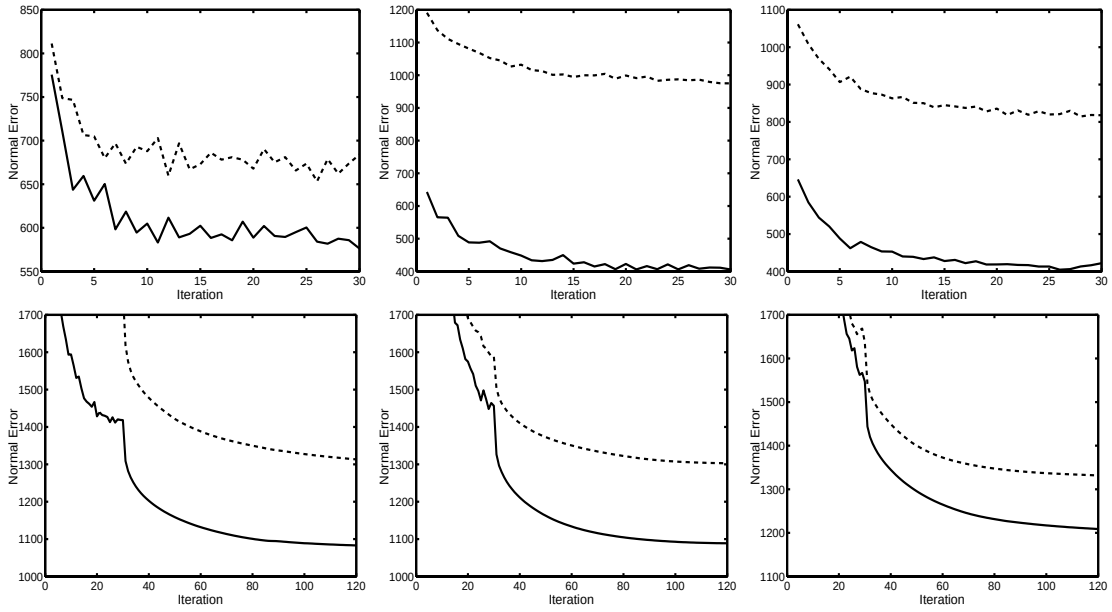

Figure 4.14: TMS surface normal error versus iteration number obtained by applying Lambertian SFS (Chapter 6) to the uncorrected (dashed) and corrected (solid) images shown in (1) Figures 4.7 and 4.8, and, (2) in Figure 4.12.

4.3.2 Real-world Data

Our experiments with real-world data are divided into three parts. We commence by illustrating the photometric correction process on a variety of objects with varying roughness. Here, we have used approximate values for the surface roughness parameter and the index of refraction parameter from tables reported in the literature for similar surfaces (Oren and Nayar, 1995; Wolff et al., 1998). We also show the results of applying Lambertian SFS to the corrected images. The second aspect of the study is to investigate in some detail the effect of the photometric correction process on the directions of the recovered surface normals. Then, we illustrate the effect of re-illuminating the recovered fields of surface

normals. Finally, we make use of ground-truth data to compare the results obtained before and after applying the correction process to the images of cylindrical surfaces.

Results of Photometric Correction

In Figure 4.15, the left-hand panel shows the raw image for a porcelain vase. In the centre panel, we show the specular reflectance reconstructed, using the Torrance-Sparrow model (Torrance and Sparrow, 1967). The right-hand panel shows the result of performing specularity removal using the method outlined in Chapter 3. This is the diffuse image used as input to the photometric correction process (Figure 4.16). However, for matte surfaces of the terra-cotta bear and the chalk owl shown in Figure 4.16, we use raw images as diffuse images. Having obtained the diffuse image for the shiny surface, and the raw images of matte surfaces, we are in the position to apply our correction process to recover the Lambertian component.

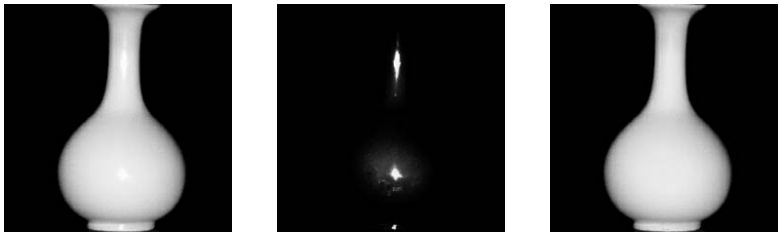

Figure 4.15: Removing (b) the specular component from (a) the raw image to obtain (c) the diffuse component using the method described in Chapter 3.

In Figure 4.16 we show the results for three objects with smooth, rough or very-rough surfaces. The objects under study are the porcelain vase (of Figure 4.15), a terra-cotta bear and the back view of a chalk owl. In the top row of the figure, we show the diffuse components I_D. Here, we use our lookup table approach for each of the three diffuse reflectance models discussed earlier. We use the Wolff model for the smooth porcelain vase and let the index of refraction be $n = 1.7$. For the very-rough chalk owl, the Oren-Nayar model is used, while for the terra-cotta bear we use the combined model. We let the surface roughness parameter be $\sigma_a = 0.12$rd for the terra-cotta bear and $\sigma_a =$

0.14rd for the chalk owl. In the second row, we show the recovered Lambertian images I_L. To display our results, we compute difference images. To emphasis the differences, we use the following procedure. Let $D_{i,j}$ be the grey-scale level at the point with pixel coordinate (i, j) in a raw image (for "matte" surfaces) or in an image to which we have applied specular subtraction (for shiny surfaces) and let $L_{i,j}$ be the grey-scale level at the same location on the recovered Lambertian image. For the purposes of visualization, we compute the quantity $DL_{i,j} = 255 - K|D_{i,j} - L_{i,j}|$ where K is a magnification coefficient. For the difference images shown in this section we use $K = 10$. This means that the darker the point in a difference image, the larger the difference between the diffuse and the Lambertian components. The third row shows the images of the differences. The maximum difference value differs between surfaces with different roughness and the absolute differences are at most 10%. Finally, in the fourth row we show transverse brightness sections across approximately cylindrical regions of the objects. Here the solid curve is the original brightness while the broken curve is the corrected brightness.

The main features to note from the difference images are as follows. First, the correction on the limbs of the objects is most marked for the porcelain vase which has smooth surface. However, the amount of limb correction for the terra-cotta bear (which has a surface with medium roughness) and the chalk owl (which has a very-rough surface) is small in both cases. Second, the amount of rough reflectance correction is greatest for the chalk owl and least marked for the porcelain vase. The differences obtained using the Wolff model are sharper than those obtained using the Oren-Nayar model. Hence, the photometric correction is steepest for the Wolff model. These experiments also show that we can apply our correction process using the combined model to different objects with different amounts of surface roughness. In particular, for surfaces with intermediate roughness the results delivered by the combined model are more reliable than those delivered by the separate rough or smooth reflectance models. For the porcelain vase, at every image location, the corrected radiance is greater than diffuse radiance. For the chalk owl the diffuse radiance is greater than the corrected radiance. Finally, for the terra-cotta bear, there are locations where the corrected Lambertian radiance is greater than the raw diffuse

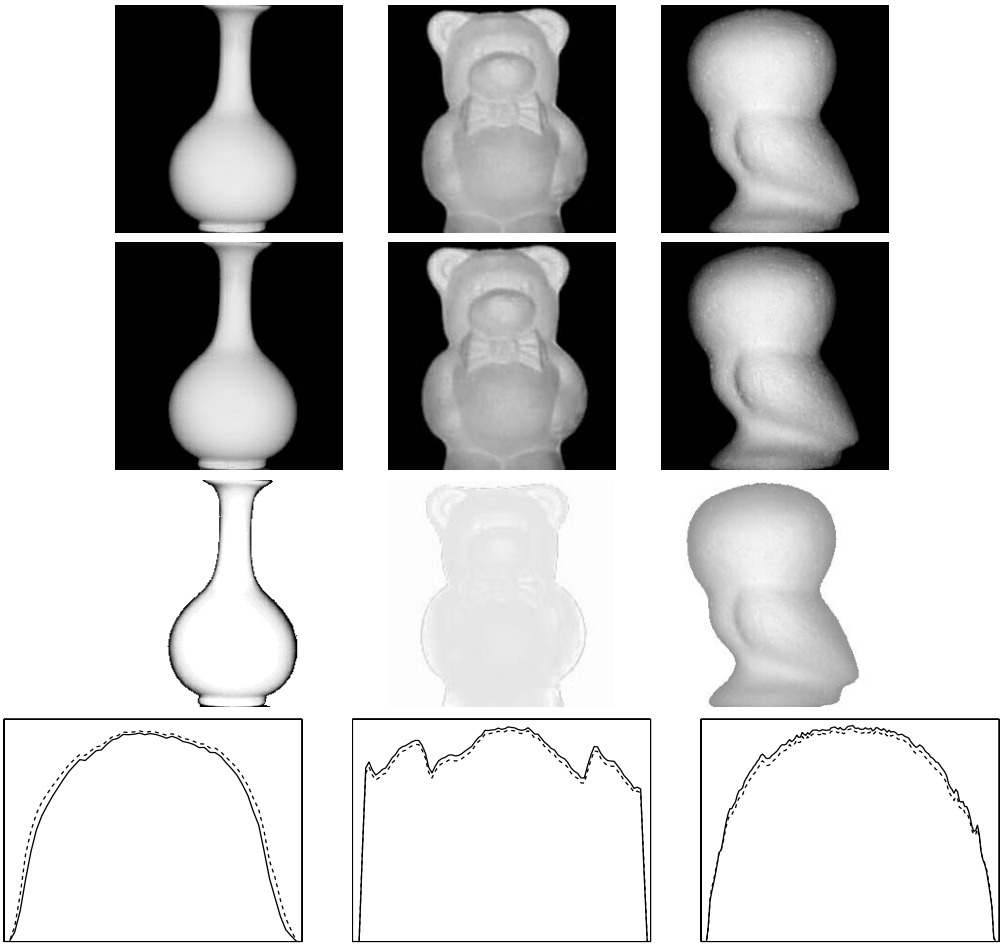

Figure 4.16: Photometric correction for (a) the porcelain vase using the Wolff model, (b) a terra-cotta bear using the combined model, and (c) a chalk owl using the Oren-Nayar model: (1) diffuse images; (2) corrected Lambertian images; (3) difference images; (4) transverse brightness sections.

radiance (i.e. only the effect of the Wolff model) and vice versa (i.e. only the effect of the Oren-Nayar model). As noted earlier, this is due to the combination of effects produced by the separate Oren-Nayar and Wolff models.

Shape-from-Shading

In this section, we turn our attention to the effects of the correction process on the surface normal directions. In Figure 4.17, we show the field of surface normals obtained by applying Lambertian SFS to the diffuse images of Figure 4.16, in the top row, and the recovered Lambertian images, in the centre row. The field of difference between these

surface normals are shown in the bottom row of the figure. Although the changes in surface normal directions are slight, they are most marked on the limbs or locations of high surface inclination. Here, the complex surface structure of the terra-cotta bear, in particular the boundaries of the arms and legs, is clearly visible. For the image of the porcelain vase, the symmetrical shape of the bulb and the neck are nicely preserved. For the images of the chalk owl, although the object has a very rough surface, the surface structure is nicely preserved. In short, the correction process results in recovering good shape information by SFS from the objects with smooth and rough surfaces.

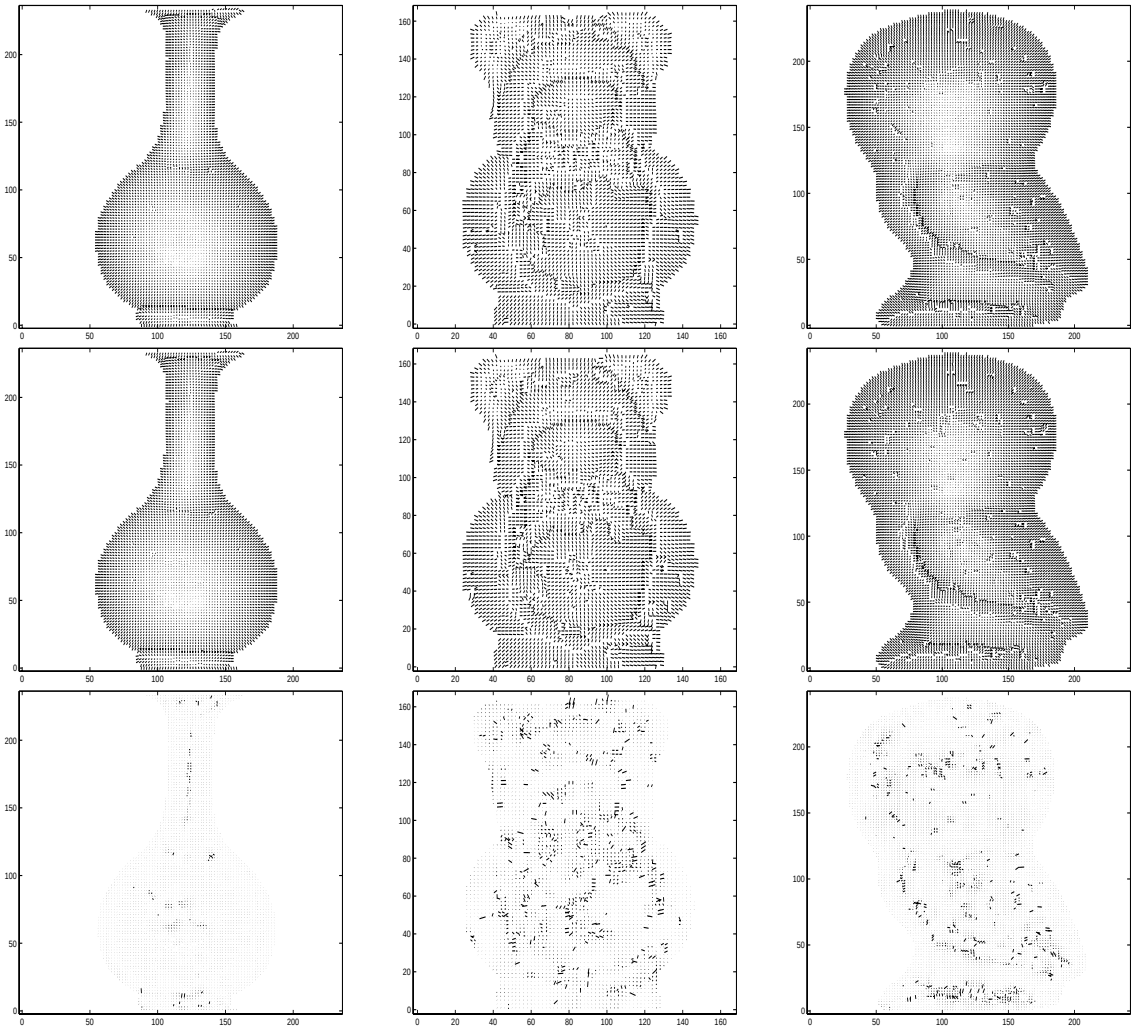

Figure 4.17: The surface normals obtained by applying Lambertian SFS (1) to the diffuse images of Figure 4.16, and (2) to their corresponding corrected Lambertian images; (3) the field of difference.

We now focus in more detail to the effects of the Lambertian correction process on the surface normal directions. In qualitative terms, when the image becomes brighter then the apex (opening) angle of the Lambertian cone becomes smaller. The reason for this is that the Lambertian SFS method constrains the surface normals to fall on a cone whose apex angle is the arc-cosine of the brightness. The axis of the cone points in the light-source direction, and since $\vec{L} = \vec{V}$, the apex angle of the cone is equal to the zenith angle. Hence the zenith angle of the surface normal decreases. When the image becomes darker, on the other hand, then the zenith angle increases. To illustrate the effect of Lambertian correction process, we have produced plots for the azimuth and zenith angles for surface normals extracted from the diffuse images (or raw images for surfaces with no specularities) and the corresponding corrected Lambertian images. In Figures 4.18-4.19, the left-hand panel is the plot for the zenith angle, while the centre panel is the plot of the azimuth angle. In both plots, the horizontal axis shows the angle extracted from the diffuse images while the vertical axis shows the angle extracted from the corrected Lambertian images. Since the normal is constrained to fall on a cone (as described above), in the zenith angle plots there is no scatter and the data points trace out a curve. The azimuth angle plots, on the other hand, do exhibit scatter. The reason for this is that the surface normals are free to rotate about the cones subject to smoothness constraints. However, in each of the azimuth plots there is a clear regression line. For the terra-cotta bear, the Lambertian zenith angle is greater than the diffuse zenith angle for small and intermediate angles. For large angles, however, the pattern is reversed. For both objects, differences in zenith angles are more marked for large zenith angles.

To investigate the stability of the surface normals under changes in azimuth angle, the right-hand panel in each figure shows the difference in azimuth angle for the diffuse and corrected images as function of the Lambertian zenith angle. From the plot for the porcelain vase (Figure 4.19.c), it is clear that the large azimuth angle differences occur at small zenith angles. Hence, these are associated with small changes in surface normal direction since the opening angle of the cone is small. From the plot for the terra-cotta bear (Figure 4.18.c), the largest differences in azimuth angle occur at intermediate zenith

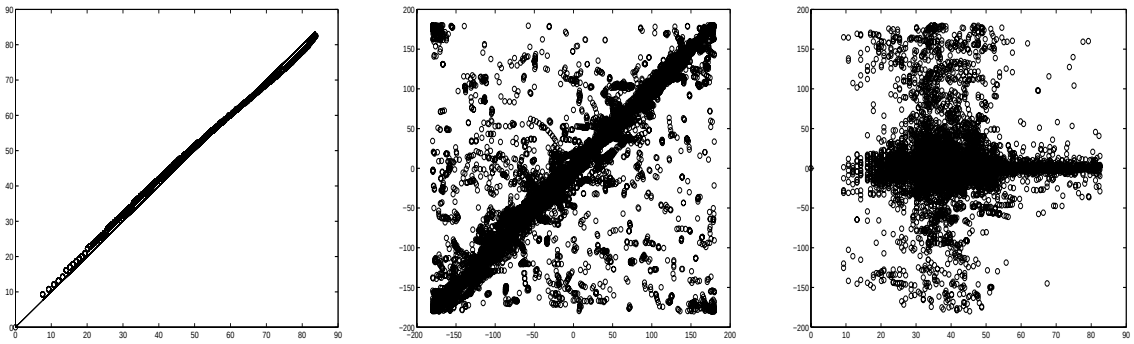

Figure 4.18: Scatter plots for the incidence (a) zenith angles and (b) azimuth angles (degrees) extracted from the recovered surface normals for the terra-cotta bear: the horizontal axis shows the angles corresponding to the diffuse image, while the vertical axis shows those for the Lambertian image (using the Oren-Nayar model); (c) the difference between the diffuse and Lambertian azimuth angles versus the Lambertian zenith angle.

angles. This means that the combined model results in significant differences in surface normal direction. Although this effect is mainly attributable to the larger differences in corrected intensity, it may also reflect problems associated with the curvature dependant smoothing employed in the SFS algorithm. This uses the principal curvature direction to adjust the azimuth angle of the surface normal. When the surface is umbilic (e.g. at spheres) or hyperbolic (e.g. at saddles) there are singularities in the field of principal curvature directions and these in turn may lead to unstable azimuth angle estimates. It is also interesting to note that SFS results in small differences in azimuth angle near object limbs. The reason for this is that the boundary condition constrains surface normal to be perpendicular to the occluding boundary.

Comparison with Ground-Truth

In this section we compare the results of the photometric correction process with ground-truth data. To do this we use cylindrical objects, which allow the evaluation of results. Hence, in a manner similar to that used in the previous chapter (Section 3.7.6), we compute the predicted surface normal and the Lambertian reflectance at each image location corresponding to the point on the cylinder fit. For our experiments we use two cylindrical surfaces with different scales of roughness. These are a strawboard cylinder and a cylin-

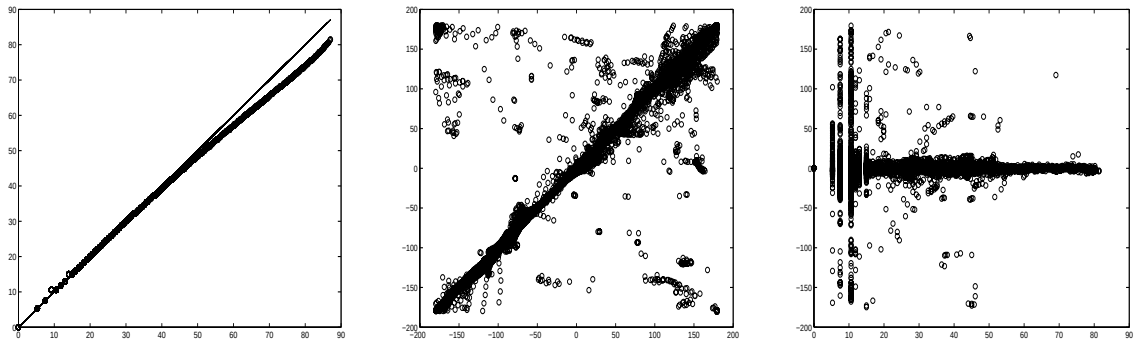

Figure 4.19: Scatter plots as described in Figure 4.18 for the porcelain vase using the Wolff model.

drical section of the porcelain vase. Where applicable we have magnification factors of $K = 5$ and $K = 1$ in the difference images. Note that for all cases, the darker a point on a difference image, the higher the difference value.

Rough Cylinder: In the top row of Figure 4.20, the left-hand panel shows the raw image of a strawboard cylinder. We show the corrected images obtained using the Oren-

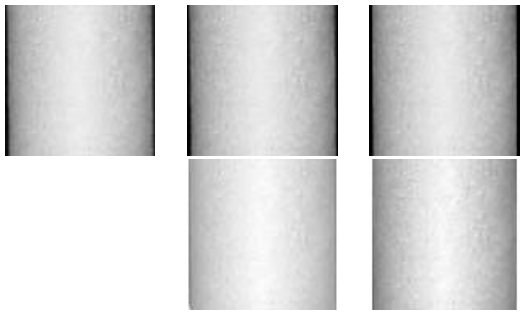

Figure 4.20: (1) Strawboard cylinder: (a) raw image; corrected images using the Oren-Nayar model (b) for $\sigma_a = 0.1$rd and (c) for $\sigma_a = 0.12$rd; (2) magnified difference between the raw and corrected images.

Figure 4.21: (a) Lambertian image for a fitted synthetic cylinder, and its absolute differences with (b) the raw and (c-d) corrected images of Figure 4.20.

Nayar model with $\sigma_a = 0.1$rd in the centre panel and with $\sigma_a = 0.12$ rd in the right-hand

panel. In the bottom row of the figure, we show the difference ($K = 5$) between the raw and the corrected images in the top row. The effect of using the larger roughness value is to increase the difference, especially near the boundary of the cylinder. Figure 4.21 illustrates the difference ($K = 1$) between the corrected Lambertian image and the predictions delivered by the fit to the cylinder geometry. In the left-hand panel of Figure 4.21 we show the Lambertian component predicted by the fit of the cylinder. The remaining panels, from left to right, show the difference between the cylinder prediction and the images in the top row of Figure 4.20, i.e. the raw images and the two corrected images. The greatest difference occurs for the raw image.

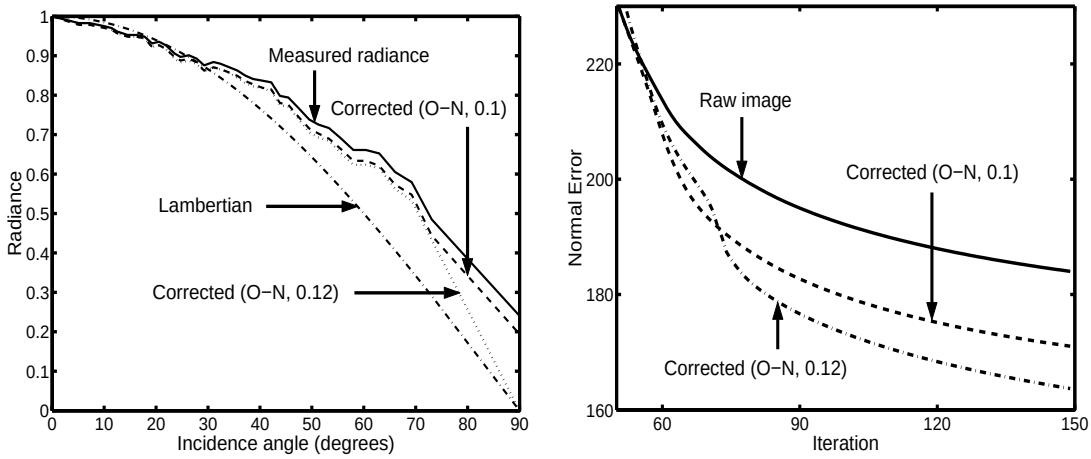

Figure 4.22: (a) Normalized raw radiance for the strawboard cylinder versus θ_i compared to the corrected radiance using the Oren-Nayar model ($\sigma_a = 0.1$ and $\sigma_a = 0.12$rd) and the Lambert's law; (b) TMS surface normal error versus iteration number obtained by applying SFS to the images of Figure 4.20.

In Figure 4.22.a, we compare the recovered Lambertian radiance using the Oren-Nayar model with the prediction of Lambert's law for a cross section on the strawboard half-cylinder. The plots show the normalized radiance as a function of the incidence angle γ computed from the fit of the cylinder. In the plot, the solid curve is the measured radiance, the dash-dot curve is the prediction of Lambert's law, the dashed curve is the recovered Lambertian radiance with $\sigma_a = 0.1$rd and the dotted curve that with $\sigma_a = 0.12$rd. Obviously, the closer the corrected Lambertian radiance to the Lambertian curve, the more successful the correction process. Here, it is the dotted curve which is the closest to

the Lambertian curve. Hence, the Oren-Nayar model with $\sigma_a = 0.12$rd results in a better correction than that with $\sigma_a = 0.1$rd. It is worth mentioning that for higher values of the roughness parameter, i.e. $\sigma_a > 0.12$rd, the difference between the corrected radiance and the ground-truth increases.

In Figure 4.22.b, we show plots of the TMS error between the recovered surface normals and those predicted by the cylinder fit, as a function of the iteration number for the SFS algorithm. The solid curve is for the raw image, the dashed curve for the corrected image obtained using the Oren-Nayar model with $\sigma_a = 0.1$rd, and the dash-dot curve that obtained with $\sigma_a = 0.12$rd. It is the dash-dot curve that corresponds to the lowest error.

Smooth Cylinder: In Figure 4.23 we turn our attention to shiny surfaces, and show some results for the porcelain vase. The left-most panel shows the original image of the whole object. In the second panel, we show the neck of the vase, i.e. the detail used in our study. The specularity is subtracted using the method described in Chapter 3. The separated specular and diffuse components are shown in the third and fourth panels respectively. The specular component is cleanly removed from the original images. In Figure 4.24, the left-most panel of the top row shows a cylindrical section of the neck of the vase from the diffuse image which is used as input to the Lambertian correction process. In the second, third and fourth panels of the top row we respectively show the corrected images obtained using the Wolff ($n = 1.7$), the Oren-Nayar ($\sigma_a = 0.1$rd) and the combined ($\sigma_a = 0.1$rd and $n = 1.7$) models. The differences ($K = 5$) between the diffuse image and each of the corrected images are shown in the corresponding panels in the bottom row of the figure. Using calibration information, we have created a synthetic cylinder of the same dimensions as the cylindrical section of the vase. In Figure 4.25, the left-most panel shows the Lambertian component obtained using the cylindrical fit to the section of the neck. In the subsequent panels of the figure, we show differences ($K = 1$) between the synthetic (ground-truth) Lambertian image, and, the separated diffuse image, the Wolff correction, the Oren-Nayar correction and the combined correction.

Before we commence our discussion of the results, we note that with glossy or shiny surfaces of this sort there are clearly potential systematic biasses associated with the spec-

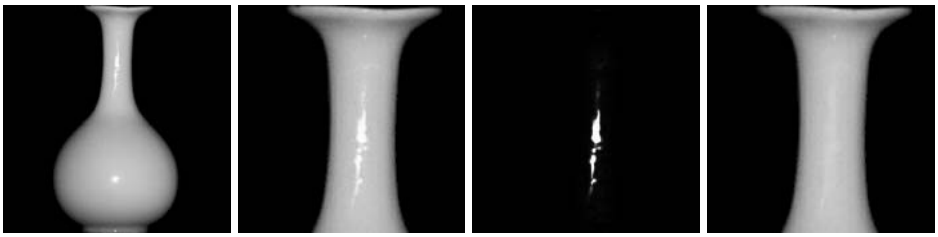

Figure 4.23: Removing specular highlights (using the method described in Chapter 3) from the raw image for the neck of a porcelain vase.

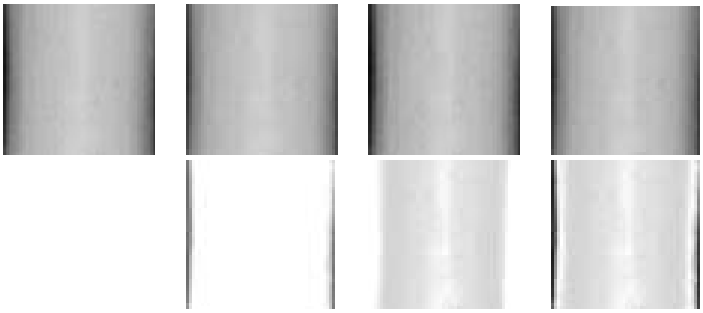

Figure 4.24: (1) Cylindrical section of the vase: (a) non-specular image, corrected images using the (b) Wolff ($n = 1.7$), (c) Oren-Nayar ($\sigma_a = 0.1$rd), and (d) combined models; (2) magnified differences between the non-specular and corrected images.

Figure 4.25: (a) Lambertian image for a fitted synthetic cylinder and its absolute differences with the (b) non-specular and (c) corrected images of Figure 4.24.

ularity subtraction process. Hence, the non-specular (diffuse) component used as input to the correction process may be in error in the proximity of the specular structure. With this caveat in mind, we note that of the three models, it is the Wolff correction that gives the poorest agreement with the predictions of the Lambertian component resulting from the cylinder fit. However, the Wolff correction process does improve the recovered surface normals. To illustrate this point, in Figure 4.26 we show TMS error versus iteration number for the SFS algorithm. Here, only the Wolff correction results in a consistently lower surface normal error curve when compared to the non-specular image. Note that since the surface of porcelain cylinder is smooth, we have used fewer iterations (90) than

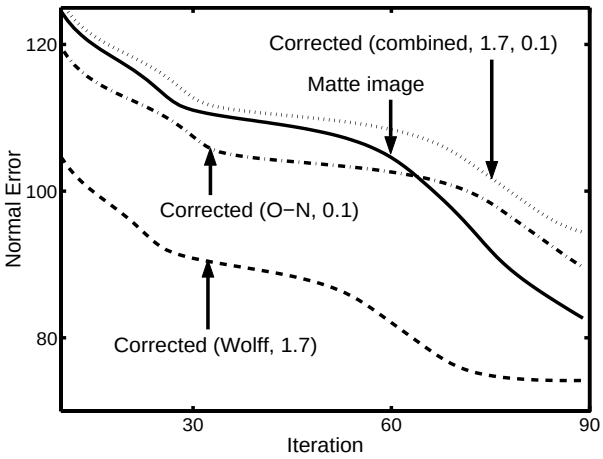

Figure 4.26: TMS surface normal error versus iteration number obtained by applying SFS to the non-specular and corrected images of the neck of the vase (Figure 4.24).

in the case of the rough surface of strawboard cylinder (150).

Re-illumination

Once accurate surface normals are in hand, we can visualize the shape of the object under study by performing re-illumination. Since at each point on the surface, the local surface normal and the viewing direction are known (and fixed), we can use the reflectance models to compute the diffuse image for any light-source direction. There is no need to use the simplified models for this purpose and the original models can be conveniently used.

Using the extracted surface normals, we hence experiment with re-illuminating surface objects from three different directions. In Figure 4.27 we show re-illuminations of the surface normals extracted from the Lambertian image of the porcelain vase in Figure 4.17 when the light-source direction is re-positioned. At each point on the surface, we compute the radiance using the three reflectance models. Here we let the roughness parameter be $\sigma_a = 0.15$rd and the index of refraction be $n = 2.0$. The panels in the left-hand column in the figure are for the Wolff model (Equation 4.10), the panels in the centre column are for the combined model (Equations 4.14 and 4.1) and the panels in the right-hand column are for the Oren-Nayar model (Equation 4.1). Here we move the direction of light-source from 0 to 90° to the image plane normal, in the horizontal direction. The rows, from top

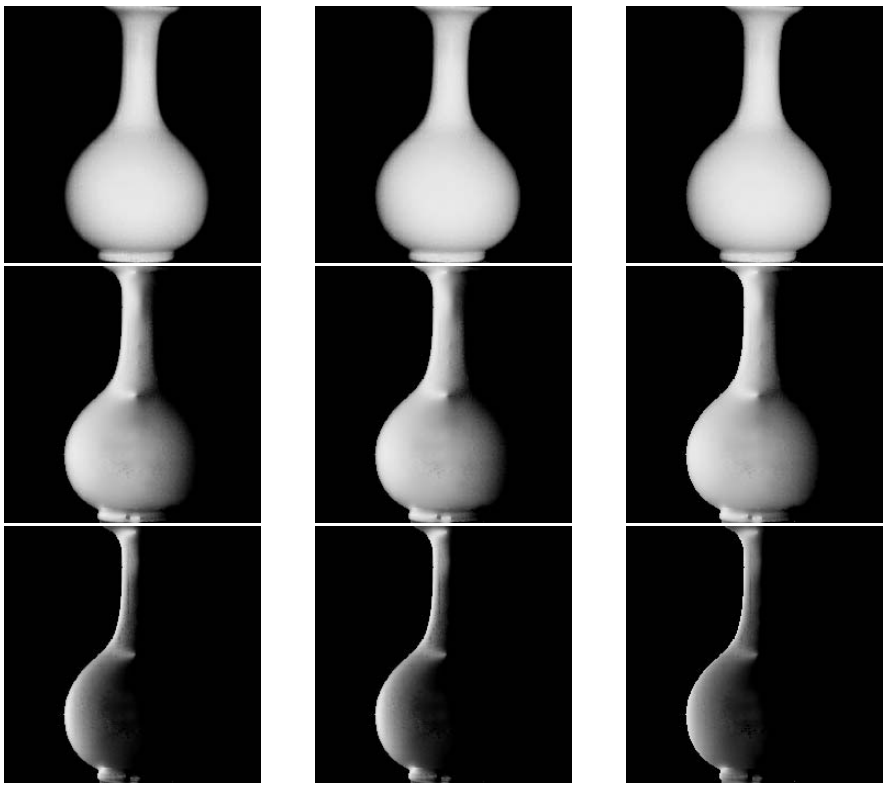

Figure 4.27: Re-illumination of the recovered surface normals (Figure 4.17) to generate diffuse images for the porcelain vase using (a) the Wolff model, (b) the combined model and (c) the Oren-Nayar model for three light-source directions: (1) $(0, 0)$, (2) $(45°, -90°)$, and (3) $(90°, -90°)$; the parameters of the models are set to $\sigma_a = 0.15$rd and $n = 2.0$.

to bottom, show the re-illuminated diffuse images for the light-source directions $(0, 0)$, $(45°, -90°)$ and $(90°, -90°)$, where the first angle is the zenith angle and the second angle is the azimuth angle of the light-source direction. The images reflect the underlying shape of the object in a consistent manner. However, there is a "pimple" on the surface. This occurs at the location of a saddle on the surface, and is due a failure of the SFS method to resolve the concave/convex ambiguity where the principal curvature direction is indeterminate. For the results shown in the first row, the images generated using the Oren-Nayar model exhibit the strongest brightness. Also, the images generated using the combined model exhibit stronger brightness than those generated using the Wolff model. However, in the bottom row, the images generated using the Oren-Nayar model exhibit the weakest brightness, while the combined model and the Wolff model result in nearly similar brightness.

4.4 Conclusions

In this chapter, we have investigated how to perform photometric correction so that a Lambertian SFS method can be applied to both shiny and rough objects. We have analyzed three diffuse reflectance models. The Wolff model is for smooth surfaces, the Oren-Nayar model is for rough surfaces and the Wolff-Nayar-Oren model combines rough and smooth scattering effects. We have derived simplified formulas for each model to perform photometric correction using table lookup. These simplified models are applicable when the angle between the light-source and viewing directions is negligible. This condition is easily achievable in practice. Under these conditions the simplified models are only dependent on the angle of incidence. Finally, we experiment with a variety of synthetic and real-world images of surfaces with different scales of roughness. Here we show the results for our photometric correction process using the three reflectance models. The results show that the combined model of Wolff-Nayar-Oren provides reliable reflectance predictions for a wide range of surfaces which can be characterized as having intermediate roughness. We have also investigated the effect of the correction process on the surface normal directions recovered using SFS. Our experiments show that different models can be used to correct the radiance provided that the appropriate roughness conditions apply, and that the correction process leads to improved surface normal estimates. Hence, we argue for the use of such diffuse reflectance models in applications including photometric stereo and surface inspection.

Chapter 5

Physics-Based Reflectance Models for Rough Surface Analysis

Our overall goal in this thesis is to develop methods which can be used to analyze shading patterns on non-Lambertian surfaces. To this end, in Chapter 3 we developed a new shape-from-shading (SFS) method using a probabilistic model in which the specular reflectance component is taken into account. We have also described how this method allows us to remove specular highlights from raw images of shiny surface. Obviously, the separated diffuse component is closer (than the raw image) to the corresponding Lambertian image. However, depending on the scales of roughness on different surfaces, the difference between the diffuse image and the ideal Lambertian image may become significant. In Chapter 4, we have moved closer to our goal by proposing a lookup table approach that can be used in conjunction with three phenomenological diffuse reflectance models to recover corrected Lambertian images from the raw (or separated) diffuse images. Unfortunately, these models are not completely based on optical physics. For instance, the Oren-Nayar model is a modified Lambertian model in which a roughness parameter is taken into account using the distribution of microfacet orientations. However, this roughness parameter does not have a purely physical meaning.

To address this omission, in this chapter, we take this study one step further by focussing on how reflectance models based on wave scattering theory can be used for rough

surface analysis. We are interested in two classes of reflectance model. For dielectric surfaces, the models investigated are derived from the Beckmann-Kirchhoff (B-K) scatter theory (Beckmann and Spizzichino, 1963), while for conducting surfaces the starting point is the Davies scattering model (Davies, 1954). We show how these two classes of model can be used to estimate surface roughness parameters for both dielectric and metallic surfaces using simple reflectance measurements. With estimates of these parameters to hand, we take our analysis one step further by demonstrating how the Vernold-Harvey modification of the Beckmann model may be used for surface analysis.

As mentioned in Chapter 2, although the B-K theory and the Davies theory are well known in physical optics they have not received significant attention in computer vision. Our aim in this chapter is to investigate whether they can be exploited for surface analysis, and focus on two applications. The first of these is how the two models can be utilized for surface roughness measurement from digital images. The second application is that of surface shape analysis. Here we show how the Vernold-Harvey modification of the B-K model may be used to recover corrected Lambertian reflectance from curved surfaces, and how this allows SFS and surface re-illumination to be applied to rough objects.

The outline of the remainder of this chapter is as follows. In Sections 5.1 and 5.2 we review prior work on physics-based modelling of rough surface reflectance. Section 5.1 focusses on the B-K scatter theory and modifications to it. Section 5.2 turns to metallic surfaces and reviews the Davies scatter theory. In Section 5.3 we describe how the models outlined in Sections 5.1 and 5.2 can be used to develop methods for estimating rough surface parameters. Section 5.4 explains how the Vernold-Harvey modification of the B-K model may be used for the purposes of Lambertian reflectance correction and surface shape analysis. Section 5.5 provides experiments on real-world data. Finally, Section 5.6 offers conclusions and directions for future investigation.

5.1 Beckmann-Kirchhoff Scatter Theory

The Beckmann-Kirchhoff (B-K) theory attempts to account for the wave-interactions of light with rough surfaces. The starting point is Kirchhoff's integral which gives the scattered light-field amplitude. Beckmann's contribution was to show how to apply the theory to rough surfaces, and how to obtain simplifications to the Kirchhoff integral under different roughness conditions. The details of the theory are comprehensively described in the monographs of Beckmann and Spizzichino (Beckmann and Spizzichino, 1963), and by Ogilvy (Ogilvy, 1991). The first aim in this study is simply to review the main results from these texts and to state the conditions under which they apply. Our second and novel contribution is to apply the results to roughness estimation and to the problem of recovering corrected Lambertian reflectance from rough surfaces.

Hence, in this section we turn our attention to the B-K theory for light scattering from rough surfaces. Based on Kirchhoff theory, Beckmann (Beckmann and Spizzichino, 1963) has found an expression for the mean scattered power (i.e. mean of the squared modulus of the scattered electric field) which is proportional to the mean optical intensity or the measured irradiance (Harvey et al., 1999). However, there are alternative schemes for obtaining these quantities from the Helmholtz-Kirchhoff integral, and so, related radiometric quantities have been computed by other authors (Stogryn, 1967; Berman and Perkins, 1985; Ogilvy, 1991; Nayar et al., 1991; Caron et al., 2002). Harvey et al. (Harvey et al., 1999) have stressed that because many diffraction-scattering theories have a paraxial limitation and there is no difference between radiance and intensity in the paraxial region (if θ is small then $\cos\theta \approx 1$), it has not been widely recognized that radiance is the fundamental quantity that is predicted accurately by scalar diffraction theory. For instance, while the experimental data collected using scattering instruments (O'Donnell and Mendez, 1987) are scattered intensity values, the predictions from the B-K theory are a better representation of the scattered radiance. Hence, the normalized form of all variants of the B-K theory presented in this chapter can be approximately used as normalized radiance quantities.

Commencing from the Helmholtz-Kirchhoff diffraction integral, and by requiring that the surface has both a Gaussian surface height distribution and a Gaussian surface correlation function, the B-K theory leads to two interesting cases. The first of these is for a surface that is smooth or of intermediate roughness. Under these conditions the following equation may be used to model the total scattered intensity:

$$I(\theta_i, \theta_r, \phi_r; \sigma, T, \lambda) = I_0 e^{-g} + \frac{\pi T^2 F^2(\theta_i, \theta_r, \phi_r) e^{-g}}{A} \sum_{n=1}^{\infty} \left[\frac{g^n}{n! n} \exp\left(\frac{-v_{xy}^2 T^2}{4n} \right) \right] \quad (5.1)$$

The model uses the geometry which is shown in Figure 5.1. Here, the incident beam

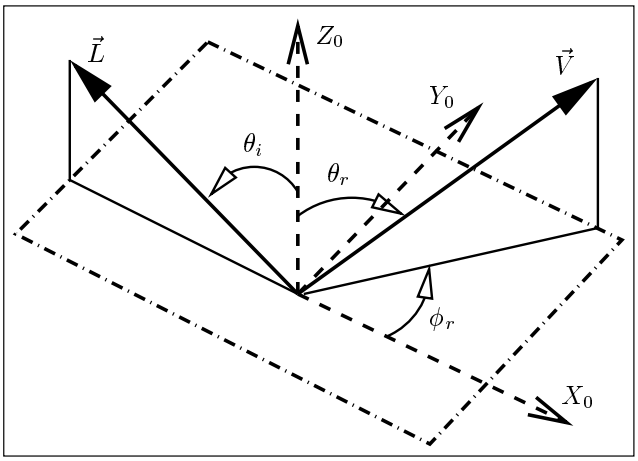

Figure 5.1: Scattering geometry on the tangent plane coordinate system: the incident light propagates in the direction $\vec{L}$, while the direction for the scattered light is $\vec{V}$.

has zenith angle θ_i and azimuth angle $\phi_i = \pi$, whereas the reflected beam has zenith angle θ_r and azimuth angle ϕ_r (on local tangent planes). From trigonometry it follows that $v_x = k(\sin\theta_i - \sin\theta_r \cos\phi_r)$, $v_y = -k(\sin\theta_r \sin\phi_r)$, $v_z = -k(\cos\theta_i + \cos\theta_r)$, $v_{xy}^2 = v_x^2 + v_y^2$, $g = \sigma^2 v_z^2$ and $k = 2\pi/\lambda$, where λ is the wavelength. The parameter σ is the root-mean-square (RMS) height deviation of the topographic surface features about the mean surface level. The height values are generally measured at equally spaced digitized data points. The correlation length T is defined as the lag-length at which either the Gaussian or the exponential correlation function drops to $1/e$ of its maximum. The

geometrical factor is given by

$$F(\theta_i, \theta_r, \phi_r) = \frac{1 + \cos\theta_i \cos\theta_r - \sin\theta_i \sin\theta_r \cos\phi_r}{\cos\theta_i(\cos\theta_i + \cos\theta_r)} \tag{5.2}$$

The quantity g has been used in the literature to divide surfaces into three broad categories. These are a) slightly-rough ($g \ll 1$), b) moderately-rough ($g \simeq 1$) and c) very-rough ($g \gg 1$) surfaces. The parameter A is the area of a plane sheet on which the scattering coefficient I_0 is defined (Beckmann and Spizzichino, 1963). The first term in Equation (5.1) determines the specular reflectance. The scattering coefficient I_0 vanishes everywhere except near the direction of specular reflection. The second term, i.e. the infinite series, determines the diffuse reflectance component. The number of terms that need to be used in the summation depends on the roughness with respect to the wavelength. Surfaces that are slightly rough have $g \ll 1$ and so the series in Equation (5.1) converges rapidly. In practice, only the first term needs to be considered and the diffuse intensity for slightly-rough surfaces becomes

$$I_d(\theta_i, \theta_r, \phi_r) \approx \frac{\pi g T^2 F^2}{A} \exp\left[-\left(g + \frac{T^2}{4} v_{xy}^2\right)\right] \tag{5.3}$$

When we consider the specular direction ($\theta_i = \theta_r$ and $\phi_r = 0$), then both v_x and v_y vanish and $F = 1$ (Figure 5.12.a). Hence the optical intensity is proportional to $\exp(-g)$. In Section 5.3, we use this equation to estimate the roughness for slightly-rough surfaces.

The second interesting case arises when the surface is very-rough when compared to the test wavelength (i.e. $g \gg 1$). Under these conditions the specular term becomes insignificant and it can be shown that the expression for the total (or diffuse) intensity is

$$I(\theta_i, \theta_r, \phi_r) \approx \frac{\pi F^2 T^2}{A v_z^2 \sigma^2} \exp\left(\frac{-v_{xy}^2 T^2}{4 v_z^2 \sigma^2}\right) \tag{5.4}$$

For different surface correlation functions there are different scattering behaviors. For instance, if the correlation function is exponential, it can be shown that the total (or dif-

fuse) intensity is

$$I(\theta_i, \theta_r, \phi_r) \approx \frac{2\pi F^2 T^2}{A v_z^2 \sigma^2 (1 + v_{xy}^2 T^2 / v_z^2 \sigma^2)^{3/2}} \tag{5.5}$$

The B-K model depends on both the incidence and the reflectance angles. However, in computer vision it is mainly the incidence angle behavior of the reflectance models that is of interest. Hence, here we derive a formulation for the specific case when the angle between the light-source direction ($\vec{L}$) and the viewing direction ($\vec{V}$) is small ($\vec{L} \approx \vec{V}$), and so $\theta_i = \theta_r = \theta$ and $\phi_r = \pi$ (Figure 5.12.d). In this case, by replacing the quantities v_z, v_{xy} and F with the corresponding functions of θ, i.e. $v_{xy} = 2k \sin\theta$, $v_z = -2k \cos\theta$ and $F = 1/\cos^2\theta$, the variant of the B-K model in Equation (5.4) reduces to

$$I(\theta) \approx \frac{T^2 \lambda^2}{16\pi A \sigma^2 \cos^6\theta} \exp\left(\frac{-T^2 \tan^2\theta}{4\sigma^2}\right) \tag{5.6}$$

For surfaces with an exponential correlation function, Equation (5.5) reduces to

$$I(\theta) \approx \frac{T^2 \lambda^2}{8\pi A \sigma^2 \cos^6\theta [1 + (T/\sigma)^2 \tan^2\theta]^{3/2}} \tag{5.7}$$

The difference between the simplified functions appearing in Equations (5.6) and (5.7) is that in the case of the Gaussian correlation function the $1/\cos^6\theta$ dependence is modified by an exponential function of $\tan^2\theta$, while in the case of the exponential correlation function the modifying factor is algebraic. The $1/\cos^6\theta$ term is singular as $\theta \to \pi/2$. In the case of Equation (5.6) this behavior is cancelled by the exponential function of $\tan^2\theta$ and the intensity tends to zero as θ approaches $\pi/2$. In the case of Equation (5.7), the intensity is singular at $\pi/2$, since the algebraic function of $\tan^2\theta$ does not cancel the singular behavior of the $1/\cos^6\theta$ term. Hence, the exponential correlation function results in a physically unrealistic behavior of Equation (5.7).

5.1.1 Modified Beckmann-Kirchhoff Model

The failure of the B-K theory to handle wide-angle scattering and large angles of incidence has been highlighted by several authors (O'Donnell and Mendez, 1987; Vernold

and Harvey, 1998). Specifically, Beckmann and Spizzichino (Beckmann and Spizzichino, 1963) have claimed that for rough reflective surfaces, it is the local surface normal direction that is of primary importance in determining the geometrical factor F. They have indicated that the local surface normal at each point on a rough surface varies as the surface is traversed, and does not coincide with the mean surface normal. Although Vernold and Harvey (Vernold and Harvey, 1998) agree with the observation that the local and mean surface normals do not coincide, they do not agree that a new obliqueness factor is needed to account for the difference in direction. Instead, they claim that it is sufficiently accurate to use the mean surface normal when modelling rough surfaces, under conditions where the wavelength of the incident light is much larger than the spatial scales under study. Hence, Vernold and Harvey (Vernold and Harvey, 1998) replace the geometrical factor (F^2) used in the B-K model with a Lambertian term that depends only on the cosine of the incidence angle ($\cos\theta_i$). This modification gives reasonable experimental agreement with scattering data for rough surfaces at large angles of incidence and large scattering angles. It should be stressed that the modification is not based on a physical model. Instead, it relies on empirical and phenomenological arguments. However, this does not diminish our pragmatic interest in the modified B-K model since it includes some effects (corresponding to non-paraxial angles) which are not included in the original B-K model. Making this replacement, the Vernold-Harvey modification to the B-K model for surfaces with a Gaussian correlation function, and simplified for $\vec{L} \approx \vec{V}$, is

$$I(\theta) \approx \frac{T^2\lambda^2}{16\pi A\sigma^2\cos\theta}\exp\left(\frac{-T^2\tan^2\theta}{4\sigma^2}\right) \tag{5.8}$$

while the corresponding result for surfaces with an exponential correlation function is

$$I(\theta) \approx \frac{T^2\lambda^2}{8\pi A\sigma^2\cos\theta[1+(T/\sigma)^2\tan^2\theta]^{3/2}} \tag{5.9}$$

Hence, the $1/\cos^6\theta$ behavior of the original B-K model (Equations 5.6 and 5.7) is replaced by a slower $1/\cos\theta$ behavior. Figures 5.2-5.3 illustrate the effect of this different

behavior for very-rough surfaces. The different curves in each panel show the normalized

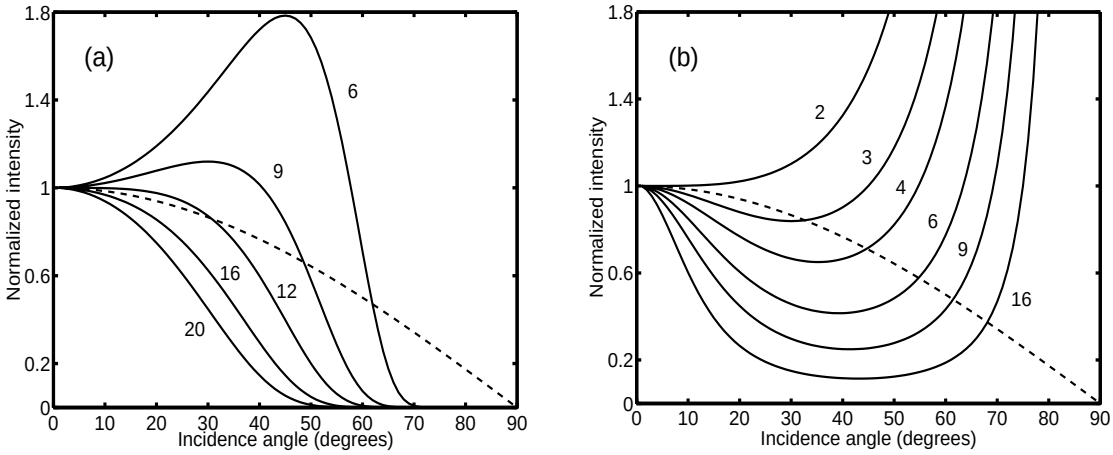

Figure 5.2: Normalized intensity using the original B-K model with (a) a Gaussian or (b) an exponential correlation function versus θ_i for several values of $\eta = (T/\sigma)^2$; Lambert's law (dashed).

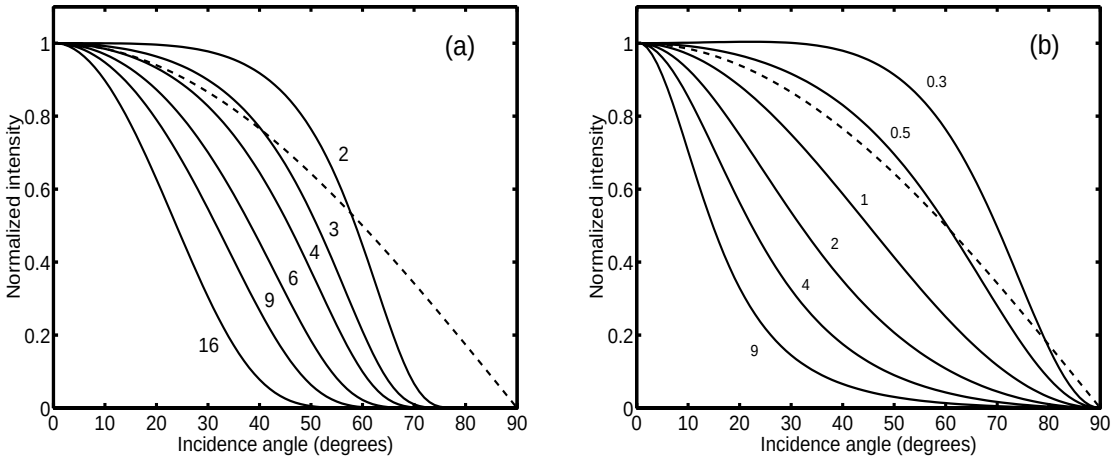

Figure 5.3: Plots (as described in Figure 5.2) for the modified B-K model.

intensity as a function of the incidence angle for different values of $\eta = (T/\sigma)^2$. The dashed curve in each panel is the prediction of Lambert's cosine law. Figure 5.2 is for the original B-K model (Equations 5.6-5.7), while Figure 5.3 is for the modified B-K model of Vernold-Harvey (Equations 5.8-5.9). In each figure, panel (a) is for the case of a Gaussian correlation function, whereas panel (b) is for the exponential correlation function. Figure 5.2 demonstrates that the original B-K model diverges significantly from the Lambertian curve for large angles of incidence. Here, in the case of the Gaussian correlation function

the normalized intensity peaks at large scatter angles for small values of η due to the $1/\cos^6\theta$ term. In the case of the exponential correlation function the behavior is singular at large scatter angles due to the discontinuity of the first derivative of this function at the origin. Figure 5.3, on the other hand, shows that the Vernold-Harvey modification to the B-K model does not exhibit these problems. From the different curves in both figures it is also clear that the larger the surface roughness σ with respect to correlation length T, the higher the normalized intensity. In particular, in Figure 5.3.a for $\eta = 3.0$ and $\eta = 2.0$, and in Figure 5.3.b for $\eta = 0.5$ and $\eta = 0.3$, the brightening effects are relatively moderate compared to the Lambertian curve.

5.1.2 Domain of Validity

Kirchhoff theory does not account for multiple scattering, which occurs in the cavities of the surface and is important for large values of the slope σ/T. When multiple scattering occurs, energy is not transferred directly from the incident wave into the scattered wave, but instead becomes "trapped" for a limited time within the hills and troughs of the rough surface (Ogilvy, 1991). Hence, Kirchhoff theory underestimates the scattered intensity. Recently, Caron et al. (Caron et al., 2002) have attempted to model this effect using the energy conservation rate. This is the ratio of the scattered energy to the incident energy, i.e. $\xi_K = E_{sc}/E_{inc}$. At normal incidence and for $g \gg 1$, when $0 \leq \theta_r \leq \pi/2$ then $\xi_K = 1 - \exp(-T^2/4\sigma^2)$. Hence, ξ_K is small if the slope is large, while it is large when the slope is low. This quantity also accounts for missing energy due to multiple scattering. This interpretation results from the fact that when the scattering process is integrated over the range $0 \leq \theta_r \leq \pi$ then $\xi_K = 1$ (Caron et al., 2002).

A second limitation of the Kirchhoff theory identified by Caron et al. is that for rough surfaces under oblique incidence there are problems due to surface self shadowing, caused by rough protrusions. To address this problem, they have introduced a limit for angles of incidence under which the predictions of the Kirchhoff model are valid. This upper limit angle is also called the angle of RMS slope and is given by $\theta_K = \pi/2 - \tan^{-1}(\sigma\sqrt{2}/T)$.

For angles of incidence greater than the angle of RMS slope θ_K, the energy scattered is significantly overestimated and the scattered intensity cannot be calculated acceptably without shadowing functions (Caron et al., 2002). Note that θ_K and ξ_K are only applicable to surfaces with a Gaussian height distribution and a Gaussian correlation function.

5.1.3 Comparison with Phenomenological Reflectance Models

In this section, we pause to compare the B-K model with the two diffuse reflectance models described in Chapter 4, i.e. the Oren-Nayar model for rough surfaces and the combined model of Wolff-Nayar-Oren for intermediate rough surfaces. In both the Oren-Nayar model and the combined model, the slope-area distribution for an isotropically rough surface is assumed to be Gaussian with a zero mean. The parameter of the distribution is the standard deviation of the cavity slope angles σ_a. Later on we will demonstrate, that of the different variants of the B-K model, it is the modified B-K model with an exponential correlation function that provides the best fit to measured reflectance data for rough and very-rough surfaces. Hence, we compare the two phenomenological reflectance models with this variant. However, it should be noted that the roughness parameter in the B-K model (σ/T) is a physics-based parameter and does not convey the same meaning as its counterpart, i.e. σ_a, in the Oren-Nayar model and in the combined model. Note that we do not perform any comparison with the Wolff model for smooth surfaces (Section 4.1.2) as there is no roughness parameter appearing in this model. However, for small values of σ_a, the combined model is equivalent to the Wolff model.

The Oren-Nayar Model

The Oren-Nayar model (Section 4.1.1) is based on geometrical optics, and hence contrasts with the wave-scattering approach of the B-K model. Recall that when $\vec{L} \approx \vec{V}$, the reflectance function of the Oren-Nayar model simplifies to

$$I(\theta) = A\cos\theta + B\sin^2\theta \tag{5.10}$$

where the parameters A and B are given in Equation (4.2). The plots in Figure 5.4 show the normalized radiance versus the incidence angle for different roughness values. In panel (a) we use $\sigma/T = 1.5$ and $\sigma_a = 0.3$rd, whereas in panel (b) we use $\sigma/T = 1.8$ and $\sigma_a = 0.6$rd. The values are chosen so that the shapes of the different distributions are roughly comparable. In the plots the solid curve is the B-K model (Equation 5.9), the dotted curve is the simplified Oren-Nayar model and the dash-dot curve is Lambert's law, i.e. $I(\theta) = \cos\theta$. The main points to note from the plots are as follows. First, the Oren-Nayar model and the B-K model are close in form for small angles. Second, the angle above which the curves diverge increases as σ_a decreases. Third, the B-K model tends to zero as θ approaches $\pi/2$, while the Oren-Nayar model does not tend to zero due to the $B\sin^2\theta$ term. Note that this limb brightening effect is a consequence of the interdependence of the coefficients A and B, and only vanishes if $\sigma_a = 0$.

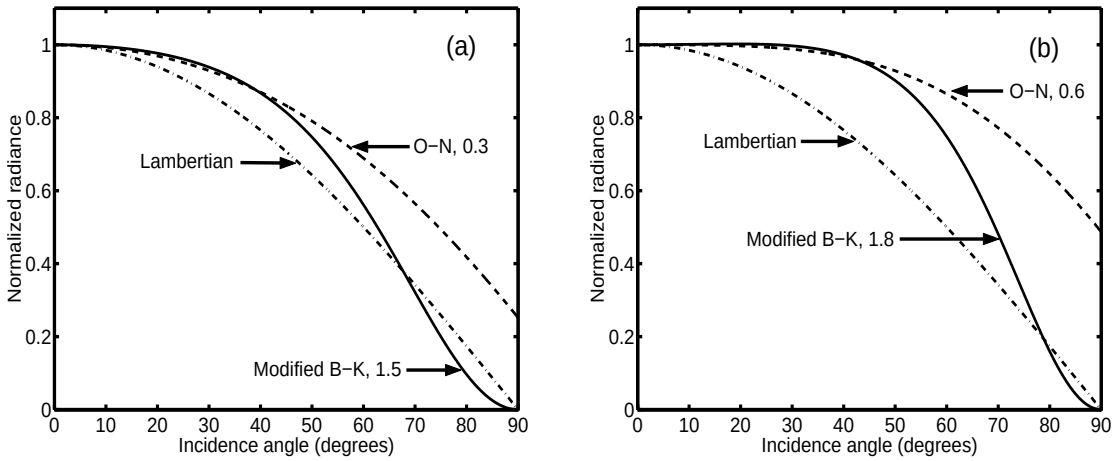

Figure 5.4: Comparing the modified B-K model (Equation 5.9) with the Oren-Nayar model (Equation 5.10): (a) $\sigma/T = 1.5$, $\sigma_a = 0.3$rd; (b) $\sigma/T = 1.8$, $\sigma_a = 0.6$rd.

The Combined Model of Wolff-Nayar-Oren

Here we compare the B-K model with the combined model (Section 4.1.3). Again, recall that for the case of $\vec{L} \approx \vec{V}$ the reflectance function of the combined model simplifies to

$$I(\theta) = A[1 - f(\theta, n)]^2 \cos\theta + B\sin^2\theta \tag{5.11}$$

The Fresnel coefficient f is given by Equations (4.11) and (4.12). This model captures a combination of effects produced by internal scattering and by the external roughness conditions. The contribution from internal scattering is modelled using the Fresnel coefficient, while the external roughness effects are modelled using the Oren-Nayar model. The plots in Figure 5.5 show the normalized radiance versus the incidence angle for different roughness values and a fixed value for the index of refraction $n = 1.7$. In panel (a) we use $\sigma/T = 1.4$ and $\sigma_a = 0.2$rd, whereas in panel (b) we use $\sigma/T = 1.2$ and $\sigma_a = 0.05$rd. The values are chosen so that the behavior of the models are roughly comparable for two surfaces of small and intermediate roughness. In the plots the solid curve is the B-K model (Equation 5.9), the dotted curve the combined model (Equation 5.11) and the dash-dot curve is Lambert's law. For panel (a) where the surface roughness is higher, the combined model and the B-K model are close in form for small angles. For panel (b) where the surface roughness is lower, the two models move closer for large angles too. Note also that for small roughness the curves for these two models are always lower than the Lambertian curve. Finally, note that the combined model in panel (b) is almost identical to the Wolff model since in this case the roughness value is very low.

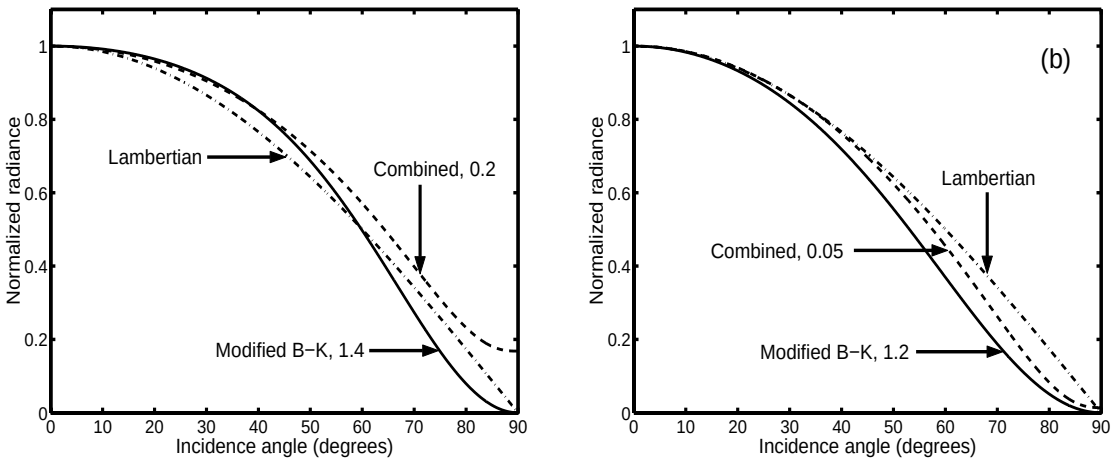

Figure 5.5: Comparing the modified B-K model (Equation 5.9) with the combined model (Equation 5.11, $n = 1.7$): (a) $\sigma/T = 1.4$, $\sigma_a = 0.2$rd; (b) $\sigma/T = 1.2$, $\sigma_a = 0.05$rd.

5.2 Davies Scatter Theory

The B-K theory applies to both dielectric and metallic surfaces. However, there are alternative models which apply for rough conducting surfaces only. The techniques reviewed in this section fall into this category and are based on the theory of reflection of electromagnetic radiation from rough surfaces, which was derived by Davies (Davies, 1954). The surface is represented by a statistical model. The four properties of the model are:

- a) the RMS roughness σ (defined as the RMS deviation of the surface from the mean surface level) is small compared with the wavelength;
- b) the surface is perfectly conducting and hence it would have a specular reflectance of unity if it were perfectly smooth;
- c) the distribution of the heights of the surface irregularities is Gaussian;
- d) the correlation function of the surface irregularities is also Gaussian (with standard deviation T).

The surface also has the statistical properties of stationarity and ergodicity with respect to position on the surface. This theory was originally developed to account for the scattering of radar waves from rough water surfaces. The model applies when the wavelength and the correlation length of the rough surface are comparable, and when this condition is satisfied, the model is also valid in the optical region. Here we consider two refinements of the Davies model which apply under different restrictions concerning the reflectance geometry. The first of these is due to Bennett and Porteus (Bennett and Porteus, 1961) and applies for normal incidence and reflectance. The second is due to Torrance (Torrance, 1964) and applies when the incident radiation is off normal.

5.2.1 Bennett-Porteus Formulation

Bennett and Porteus (Bennett and Porteus, 1961) described a technique for determining the surface roughness of opaque materials using an optical method. They obtain the fol-

lowing expression for the reflection ratio from a conducting surface at normal incidence

$$\frac{R}{R_0} = \exp\left[-\left(\frac{4\pi\sigma}{\lambda}\right)^2\right] + \frac{2^5\pi^4}{m^2}\left(\frac{\sigma}{\lambda}\right)^4 (\Delta\theta)^2 \tag{5.12}$$

where σ is the RMS roughness, m is the RMS slope of the surface profile (also $m = \sigma\sqrt{2}/T$ if T is the correlation length), λ is the wavelength of incident radiation, $\Delta\theta$ is the acceptance angle, R is the measured reflectance and R_0 is the specular reflectance of a perfectly smooth surface of the same material. The exponential term appearing in Equation (5.12) represents the specular reflectance, while the second term is due to diffusely reflected energy. Bennett and Porteus have drawn attention to the strong inverse dependency of the diffuse component on wavelength. In particular, they have noted that, at wavelengths which are long compared to the roughness, the diffuse contribution is negligible. Hence, by measuring the reflectance at two wavelengths (one of which is long enough so that the effect of the slope may be neglected) it should be possible to determine both the surface roughness and the RMS slope of the surface irregularities. When the wavelength is long enough so that the diffuse reflectance may be neglected, Equation (5.12) may be written as

$$\frac{R}{R_0} = \exp\left[-\left(\frac{4\pi\sigma}{\lambda}\right)^2\right] \tag{5.13}$$

This equation may be conveniently used to calculate the value of the RMS roughness from experimentally measured values of the reflectance at normal incidence. The estimates of the roughness obtained in this way are good approximations even if the contribution from the diffuse reflectance is not negligible. In the limit of small roughness (i.e. $\sigma \ll \lambda$) the diffuse reflectance is related to the RMS height through the total integrated scattering (TIS) ratio

$$TIS = \frac{R_d}{R_0} = \left(\frac{4\pi\sigma}{\lambda}\right)^2 \tag{5.14}$$

where R_d is the total diffuse reflectance, i.e. $R_d = R_0 - R$. If the TIS ratio (i.e. R_d/R_0) is measured at a fixed wavelength, then the surface roughness σ is directly proportional to $(TIS)^{1/2}$. It is interesting to note that Equation (5.14), on which TIS measurements are

based, can be obtained not only from Davies theory but also either from Kirchhoff theory (for surfaces with a Gaussian height distribution), or from perturbation theory when it is assumed that $\theta_i \approx \theta_r$ (Stover, 1995). The latter assumption is valid for near-normal incidence. In both approaches the condition $\sigma \ll \lambda$ is necessary. This method for obtaining a value of RMS roughness from a measurement of TIS forms the basis for a U.S. standard of effective surface roughness (Bennett and Mattsson, 1999).

5.2.2 Torrance Formulation

Equation (5.12) applies for normal incidence and reflectance, but when the incident radiation is off normal, the equation derived by Torrance (Torrance, 1964) for the reflectance ratio applies. In this case

$$\frac{R}{R_0} = \exp\left[-\left(\frac{4\pi\sigma\cos\theta}{\lambda}\right)^2\right] + \frac{2^5\pi^3}{m^2}\left(\frac{\sigma}{\lambda}\right)^4 \cos^3\theta(\Delta\omega) \tag{5.15}$$

where $\Delta\omega$ is the acceptance solid angle of the instrument. Equation (5.15) reduces to Equation (5.12) when $\theta = 0$ and $\Delta\omega = \pi(\Delta\theta)^2$. Figure 5.6 shows the relative intensity (R/R_0) against the grazing angle ($\pi/2 - \theta$) using the Torrance formula. The different curves are for different values of the scale of roughness (σ/λ). As the scale of roughness increases, then so the large grazing angle tail of the distribution increases.

In Figures 5.7-5.8 we compare the Bennett-Porteus and the Torrance formulas. Figure 5.7 is for the Bennett-Porteus formula (Equation 5.12) while Figure 5.8 is for the Torrance formula (Equation 5.15). In both figures, in panel (a), the different curves are the relative scattered intensity (R/R_0) versus the RMS slope m for different values of σ/λ. In panel (b), the different curves show the relative scattered intensity as a function of the wavelength λ for different values of m and a fixed value of σ. The dashed curves in both Figures 5.7.b and 5.8.b are for long wavelengths, at which the diffuse reflectance contribution is negligible. Here, as the RMS slope m increases the curves move closer to the dashed curve. Also, in Figures 5.7.a and 5.8.a, for $m > 0.15\mu m$, as the scale of roughness σ/λ increases from 0.05 to 0.2, the relative intensity decreases. This is because in both

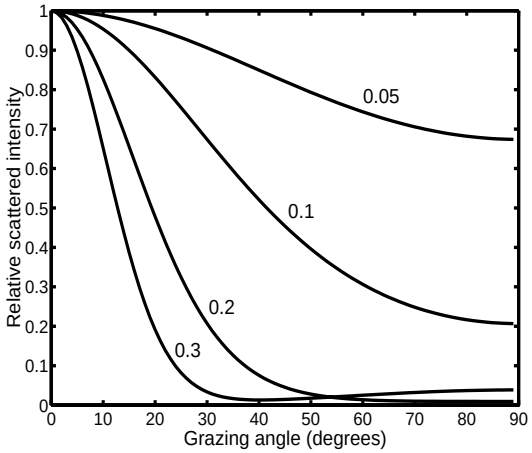

Figure 5.6: The Torrance formula: relative intensity versus grazing angle ($\pi/2 - \theta_i$) for $m = 0.9\mu m$, $\Delta\omega = 0.0039$sr, and $\sigma/\lambda = 0.05, 0.1, 0.2, 0.3$.

Equation (5.12) and Equation (5.15), for these values of σ/λ, the second term is negligible compared to the first term. However, when σ/λ increases from 0.2 to 0.3, the relative intensity increases too. The reason for this is that for such values of the roughness scale (e.g. $\sigma/\lambda \geq 0.3$) the second term, which is proportional to $(\sigma/\lambda)^4$, becomes large enough to compensate for the decrease in the exponential terms. Note also that the behavior of the relative intensity plots changes again for large values of m.

5.3 Surface Characteristic Estimation

Previous methods for determining surface characteristics, and in particular roughness, work at long wavelengths and use special purpose instruments to make reflectance measurements. However, we perform our experiments in the visible region of the spectrum and make use of an ordinary digital camera. To avoid possible noise, saturation or digitization artifacts, we compute the mean-intensity value by averaging the measured reflectance values over a neighborhood of points in the image. This technique agrees with the definition of mean optical intensity from the Beckmann scattering theory (Beckmann, 1967). Furthermore, we need to know neither the strength of the light-source nor the surface albedo. The reason for this is that the techniques discussed in this chapter use the

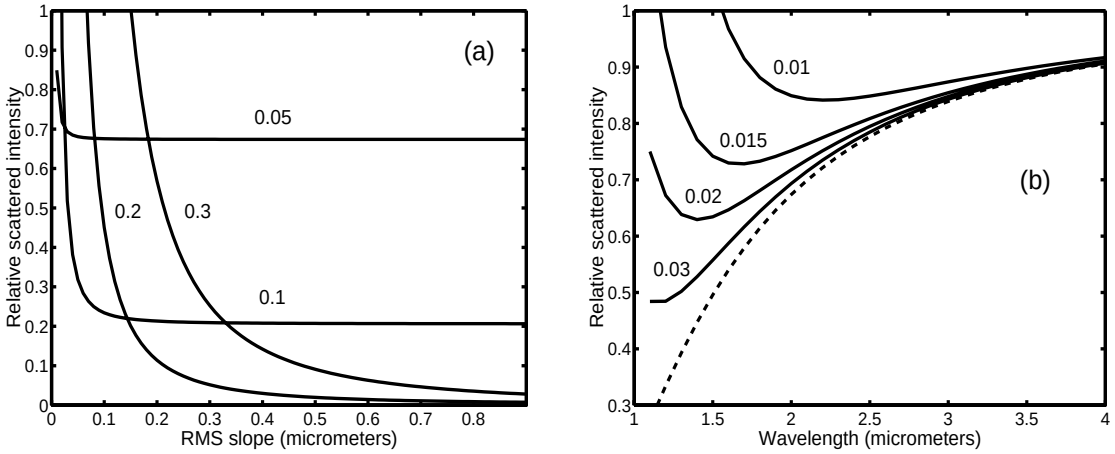

Figure 5.7: The Bennett-Porteus formula: relative intensity (for $\Delta\theta = 0.03$sr) versus (a) RMS slope (μm) for $\sigma/\lambda = 0.05, 0.1, 0.2, 0.3$, (b) versus wavelength (μm) for $\sigma = 0.1\mu m$ and $m = 0.01, 0.015, 0.02, 0.03\mu m$.

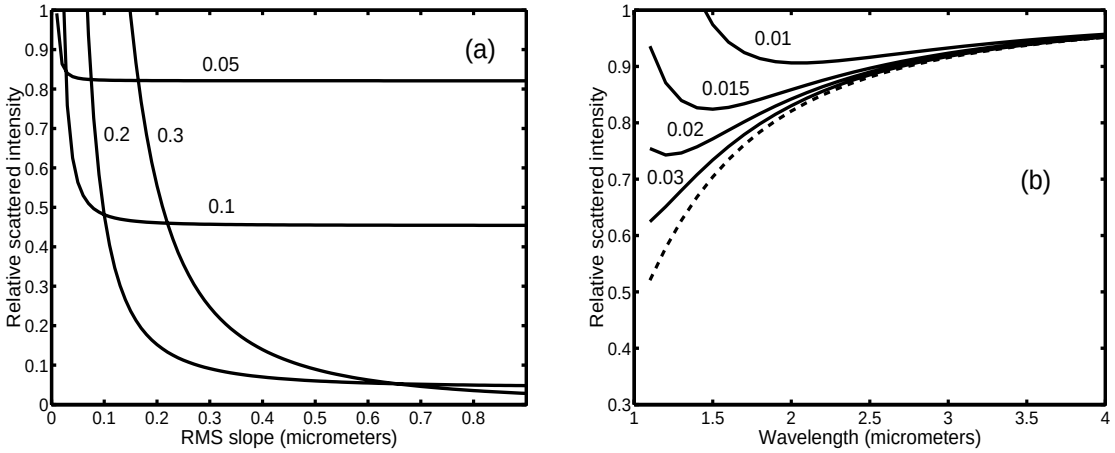

Figure 5.8: The Torrance formula: relative intensity (for $\theta = \pi/4$rd and $\Delta\omega = 0.0039$sr) versus (a) RMS slope (μm) for $\sigma/\lambda = 0.05, 0.1, 0.2, 0.3$, (b) versus wavelength (μm) for $\sigma = 0.1\mu m$ and $m = 0.01, 0.015, 0.02, 0.03\mu m$.

ratio of two intensities, and not the absolute intensity values. For very-rough dielectrics, our main contribution is as follows. We suggest a method for estimating a measure of the surface slope σ/T. We do this using intensity values measured at two different angles of illumination incidence. In practice, we capture the first image with the planar surface perpendicular to the viewing direction. The second image is captured with the camera and the object fixed, but the light-source direction is moved by a predetermined amount (Figure 5.12.b). We then compute the mean-intensity ratio, which we define as the ratio

of the mean-intensity values from the two images. From the mean-intensity ratio, we can compute σ/T using the B-K model.

When the lateral correlation length is larger than one wavelength and the surface slope is not large, then measurements of diffuse scattering are found to agree well with the B-K model, provided that the angle of incidence is not too large (O'Donnell and Mendez, 1987). When the incidence angle is large, we use the modification of the B-K model developed by Vernold and Harvey (Vernold and Harvey, 1998).

Having reviewed the various approaches available for modelling reflectance from rough surfaces, in this section we turn to the first of our surface analysis applications. In particular, we present some techniques for estimating the different physical parameters for rough surfaces described earlier in this chapter, i.e. the RMS roughness σ, the correlation length T and the RMS slope m. For surfaces with a Gaussian correlation function, if two of these quantities are known, then the third one may be obtained using the formula $m = \sigma\sqrt{2}/T$. Note that the quantity $\sigma_g = \sigma\sqrt{2}/T$ has also been referred to as the RMS gradient (Ogilvy, 1991). However, for very-rough surfaces we can only estimate the ratio of the RMS roughness to the correlation length σ/T which is repeatedly used in Equations (5.4) and (5.5). It is worth mentioning that He et al. (He et al., 1991) who used the Kirchhoff theory, also stated that the surface RMS slope is proportional to the ratio σ/T. The estimated ratio σ/T allows us to compute the angle of RMS slope $\theta_K = \pi/2 - \tan^{-1}(\sigma\sqrt{2}/T)$ of Caron et al (Caron et al., 2002). For incidence angles smaller than θ_K, the reflectance predictions delivered by the Kirchhoff model are reliable. For conducting surfaces, although we have presented a method for estimating the RMS slope m in Section 5.2.1, we do not intend to experiment at long wavelengths. Therefore, we consider only three groups of surfaces. These are dielectrics with small or large roughness, and metallic surfaces.

5.3.1 Dielectric Surfaces with Small Roughness

The first case that we consider is to estimate σ and T for slightly-rough dielectric surfaces.

Estimating RMS Roughness σ

Here, we describe a method for estimating the surface roughness parameter using specular reflectance. When the surface is smooth enough to produce a well defined specular beam, the RMS roughness can be determined from the relative intensity of the specular beam (Bennett and Porteus, 1961; Marx and Vorburger, 1990). According to the B-K theory (Beckmann and Spizzichino, 1963), the relative intensity of the specular beam is given by

$$i_{spec} = \exp\left[-\left(\frac{4\pi\sigma\cos\theta}{\lambda}\right)^2\right] \tag{5.16}$$

where θ is the incidence angle. In addition to the assumption of a Gaussian height distribution, the smooth surface assumption ($\sigma \ll \lambda$) is also made, i.e. only light scattered close to the specular beam is considered. In virtually all cases, the latter assumption is justified because most scattering occurs within a cone with an apex angle of 20° or less about the specular reflection (Stover, 1995). The expression given in Equation (5.16) uses essentially the same three characteristic quantities (σ, λ, θ) as the well known Rayleigh criterion with which it is consistent (Beckmann, 1967). Figure 5.9.a is a plot of the relative intensity of the specular beam given by Equation (5.16) for different values of the scale of roughness (σ/λ). The vertical axis is for the relative scattered intensity in the specular direction (i_{spec}), whereas the horizontal axis is for the grazing angle.

One method is to measure all the scattered light (diffuse component) from the surface and then calculate the total intensity from the sum of the specular and diffuse intensity components. However, we determine an experimental value for the relative intensity i_{spec} by measuring both the total peak intensity $\hat{I}_{tot}$ at the center of the specular highlight and the average diffuse intensity $\hat{I}_{dif}$ in the neighborhood of the specular lobe. The relative intensity can then be approximated using the formula

$$i_{spec} = 1 - \frac{\hat{I}_{dif}}{\hat{I}_{tot}} \tag{5.17}$$

For the specular direction the angle of incidence is equal to the angle of reflectance.

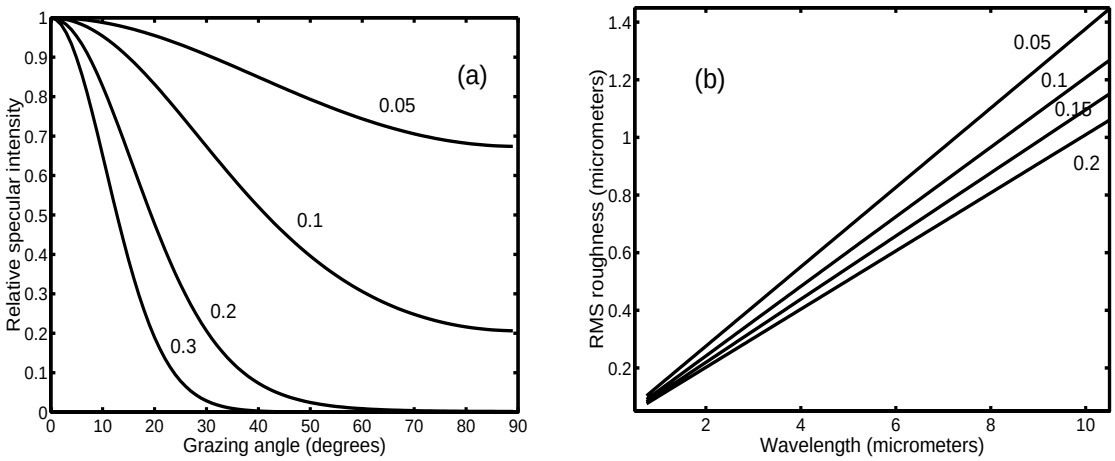

Figure 5.9: The B-K model for the specular direction: (a) relative intensity (i_{spec}) versus grazing angle for $\sigma/\lambda = 0.05, 0.1, 0.2, 0.3$; (b) σ versus λ for $i_{spec} = 0.05, 0.1, 0.15, 0.2$.

So, the angle $\theta = \Psi/2$ can be found by measuring the angle Ψ between the light-source and viewing directions (Figure 5.12.a). Also, the wavelength λ of the light-source is assumed known. Hence, the RMS roughness is computed using the formula

$$\sigma = \frac{\lambda}{4\pi \cos\theta} \left[\ln\left(\frac{1}{i_{spec}}\right) \right]^{1/2} \tag{5.18}$$

In Figure 5.9.b, we show plots of the RMS roughness computed using Equation (5.18). The vertical axis is the RMS roughness σ, while the horizontal axis is the wavelength λ. The different curves are for different values of relative intensity i_{spec}.

Estimating Correlation Length T

With a value of the RMS roughness σ to hand, then we can proceed to measure the correlation length T. To do this we measure two values of reflectance I_1 and I_2 at a single point on the surface for two different angles of incidence, i.e. $\theta_i = \theta_1$ and $\theta_i = \theta_2$ (Figure 5.12.b). Here, the planar surface is perpendicular to the viewing direction for both measurements, i.e. $\theta_r = 0$. The intensity measurement are averaged over a neighborhood so that the effects of rough surface texture are smoothed away. Although the use of two different wavelengths is also possible, here we use only one wavelength. Under

such illumination conditions and from Section 5.1 we can write $v_{xy}(\theta_i) = (2\pi/\lambda)\sin\theta_i$, $v_z(\theta_i) = (2\pi/\lambda)(1+\cos\theta_i)$ and $g(\theta_i) = \sigma^2 v_z^2(\theta_i)$. Note that when experimenting with intermediate and large incidence angles, we should use the Vernold-Harvey modification and replace F^2 with $\cos\theta_i$ in all equations. Using Equation (5.3), the ratio I_1/I_2 is

$$\frac{I_1}{I_2} = \frac{g(\theta_1)F^2(\theta_1)}{g(\theta_2)F^2(\theta_2)} \exp\left\{ g(\theta_2) - g(\theta_1) + \frac{T^2}{4}\left[v_{xy}^2(\theta_2) - v_{xy}^2(\theta_1)\right]\right\} \tag{5.19}$$

Since all of the parameters appearing in the equation except T are known, we can estimate the correlation length T using the equation

$$T = 2\left[\frac{g(\theta_1) - g(\theta_2) + \ln\kappa}{v_{xy}^2(\theta_2) - v_{xy}^2(\theta_1)}\right]^{1/2} \tag{5.20}$$

where $\kappa = [g(\theta_2)F^2(\theta_2)I_1]/[g(\theta_1)F^2(\theta_1)I_2]$.

5.3.2 Dielectric Surfaces with Large Roughness

For very-rough dielectrics, the B-K model can be used to estimate the ratio σ/T from a pair of reflectance measurements (I_1, I_2) collected at two different incidence angles θ_1 and θ_2 (Figure 5.12.b). From Equation (5.4) for the B-K model with a Gaussian correlation function ($T = T_G$), the ratio I_1/I_2 is given by

$$\frac{I_1}{I_2} = \frac{F^2(\theta_1)v_z^2(\theta_2)}{F^2(\theta_2)v_z^2(\theta_1)} \exp\left\{ \frac{{T_G}^2}{4\sigma^2}\left[\frac{v_{xy}^2(\theta_2)}{v_z^2(\theta_2)} - \frac{v_{xy}^2(\theta_1)}{v_z^2(\theta_1)}\right]\right\} \tag{5.21}$$

Hence, the ratio σ/T_G, which is a measure of the surface slope, is given by

$$\frac{\sigma}{T_G} = \frac{1}{2}\left\{\frac{1}{\ln\zeta}\left[\frac{v_{xy}^2(\theta_2)}{v_z^2(\theta_2)} - \frac{v_{xy}^2(\theta_1)}{v_z^2(\theta_1)}\right]\right\}^{1/2} \tag{5.22}$$

where $\zeta = [F^2(\theta_2)v_z^2(\theta_1)I_1]/[F^2(\theta_1)v_z^2(\theta_2)I_2]$. Again, when experimenting with relatively large angles of incidence, the Vernold-Harvey modification should be used, i.e. F^2 should be replaced by $\cos\theta_i$. This procedure may be repeated for surfaces with an

exponential correlation function ($T = T_E$) using Equation (5.5). The intensity ratio is

$$\frac{I_1}{I_2} = \frac{F^2(\theta_1)v_z^2(\theta_2)}{F^2(\theta_2)v_z^2(\theta_1)} \left\{ \frac{1 + ({T_E}^2/\sigma^2)[v_{xy}^2(\theta_2)/v_z^2(\theta_2)]}{1 + ({T_E}^2/\sigma^2)[v_{xy}^2(\theta_1)/v_z^2(\theta_1)]} \right\}^{3/2} \tag{5.23}$$

Hence, the ratio σ/T_E is computed using the formula

$$\frac{\sigma}{T_E} = \left\{ \frac{1}{(\zeta^{\frac{2}{3}} - 1)} \left[\frac{v_{xy}^2(\theta_2)}{v_z^2(\theta_2)} - (\zeta^{\frac{2}{3}}) \frac{v_{xy}^2(\theta_1)}{v_z^2(\theta_1)} \right] \right\}^{1/2} \tag{5.24}$$

In Equations (5.21) to (5.24), T_G and T_E denote the correlation lengths for the Gaussian and the exponential correlation functions, respectively. From the estimated ratio σ/T_G we can easily find the two quantities θ_K and ξ_K described in Section 5.1.2.

5.3.3 Metallic Surfaces with Small Roughness

We now turn our attention to estimating the RMS roughness for metallic surfaces. We confine our attention to surfaces with a Gaussian correlation function. Although this may not be accurate if the surface is better described by an exponential correlation function, it is the approach followed by Davies (Davies, 1954). Our aim is to estimate the RMS roughness using Equation (5.13). To make use of this formula, the illuminant direction must be normal to the surface (Figure 5.12.c), and we require measurements of the total diffuse reflectance R_d and the total reflectance R_0. We estimate the total diffuse reflectance ($\hat{R}_d$) by averaging the image intensity over a neighborhood around the specular lobe. We repeat this procedure for a neighborhood around points with maximum specular intensity to find an estimate for the mean specular reflectance ($\hat{R}_0$), which we use as an estimate of the total reflectance. Hence, the TIS ratio is given by $TIS = \hat{R}_d/\hat{R}_0$. The RMS roughness may then be computed from Equation (5.13) using the formula

$$\sigma = \frac{\lambda}{4\pi}[-\ln(1 - TIS)]^{1/2} \tag{5.25}$$

When the wavelength is long relative to the roughness ($\lambda \gg \sigma$), RMS roughness

can be found using Equation (5.14), which is a simplified form of Equation (5.13). The approximate roughness is given by

$$\sigma = \frac{\lambda}{4\pi}(TIS)^{1/2} \tag{5.26}$$

Once an estimate of the RMS roughness σ is to hand, Equation (5.15) may be used to find an estimate of the RMS slope m. In doing this, we measure the relative diffuse reflectance $\hat{R}_d/\hat{R}_0$ for an off normal incidence angle θ while the viewing direction is aligned in the specular direction (Figure 5.12.a). This allows us to find the correlation length T using the formula $T = \sigma\sqrt{2}/m$.

5.4 Exploiting the Beckmann Model

Once the physical surface parameters have been estimated using the methods outlined in Section 5.3, then the B-K model may be used to model the reflectance from curved rough surfaces. Since the angles encountered are likely to be large, it is the Vernold-Harvey modification to which we turn our attention. Although the model is intractable for the purposes of SFS, it may be used to recover a corrected Lambertian image in the manner already described in Chapter 4. Hence, surface normal recovery can be effected by applying a standard Lambertian SFS method to the corrected image. These surface normals may be used to perform novel view synthesis of rough surfaces through object re-illumination from different angles.

5.4.1 Photometric Correction

We are interested in using the B-K theory to recover corrected Lambertian image under conditions in which $\vec{L} \approx \vec{V}$ (Figure 5.12.d). To estimate the surface roughness parameters in Section 5.3, we could use the original B-K model since the surface samples were globally flat and it was sufficient to confine our experiments to small angles of incidence and scatter. However, for curved surface analysis, we can not control the geometry in this

way. Hence, we need to make use of the large angle modification of Vernold and Harvey. Moreover, we must turn our attention to two distinct roughness conditions where different variants of the modified B-K model apply. The first of these is for very-rough surfaces, where an estimate of the surface slope parameter $m = \sigma/T$ is required. The second case is for slightly-rough surfaces where separate estimates of the surface roughness σ and the correlation length T are required.

Very-Rough Surfaces

For very-rough surfaces we confine our attention to the exponential correlation function since this gives a better fit to measured surface data (Bennett and Mattsson, 1999; Ogilvy, 1991). The Vernold-Harvey modification which applies to this case under conditions of $\vec{L} \approx \vec{V}$ (Figure 5.12.d) is given in Equation (5.9). We can rewrite this equation to make the role of the RMS slope m more explicit. The result is

$$I(\theta_i) \approx \frac{\lambda^2}{8\pi A m^2 \cos\theta_i [1 + (\tan^2\theta_i / m^2)]^{3/2}} \tag{5.27}$$

Note that the normalized intensity does not depend on λ or A. In Figure 5.10.a, we show the behavior of the modified B-K model given by Equation (5.27). The different

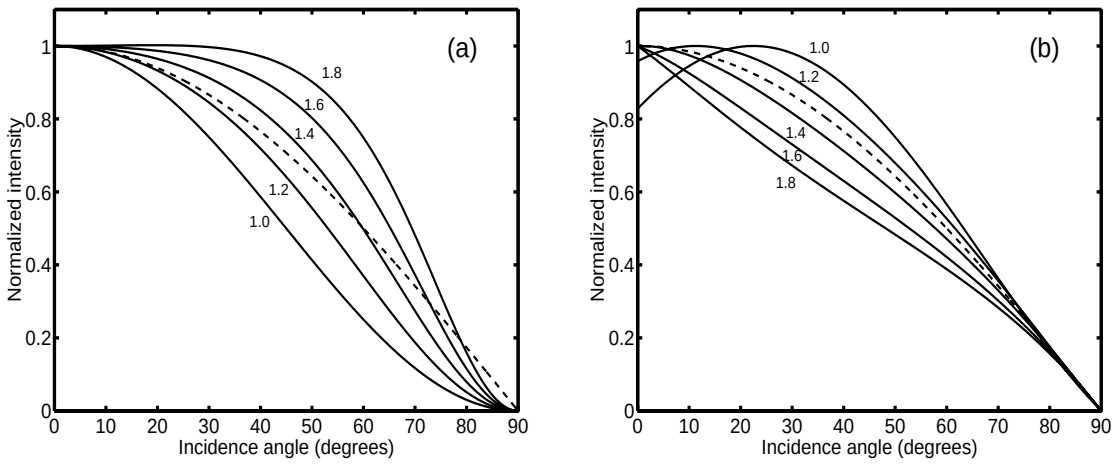

Figure 5.10: Normalized intensity using the modified B-K model (Equation 5.5) versus incidence angle for $m = 1.0, 1.2, .., 1.8$ compared to the Lambertian model (dashed) for two illumination conditions: (a) $\theta_i = \theta_r$ and $\phi_r = \pi$, (b) $\theta_i = \pi/2 - \theta_r$ and $\phi_r = 0$.

curves in each panel show the normalized intensity I as a function of the incidence angle θ_i for different values of surface slope $m = \sigma/T$. The dashed curve in the figure is the prediction of Lambert's cosine law. From the different curves it is clear that the larger the surface slope, (i.e. the surface roughness σ with respect to the correlation length T), the higher the intensity. In particular, for $m = 1.6$ and $m = 1.8$, the brightening effects are significant compared to the Lambertian reflectance curve. In Figure 5.10.b, we show the intensity curves obtained using Equation (5.5) when the light-source direction is perpendicular to the viewing direction, i.e. $\theta_i = \pi/2 - \theta_r$ and $\phi_r = 0$. By contrast with Figure 5.10.a, as the surface slope decreases from $m = 1.8$ to $m = 1.0$, then so the curves move higher. Also, the maximum intensity moves away from $\theta_i = 0$ and occurs when the incidence angle is close to $\theta_i = 25^\circ$ (for $m = 1.0$). This behavior also means that the reflectance function is not injective, and so, not invertible.

Slightly-Rough Surfaces

We now turn our attention to slightly-rough surfaces. Hence, we commence from the simplified form of the B-K model given in Equation (5.3). To apply the Vernold-Harvey modification to this variant, we replace the term F^2 with $\cos\theta_i$. Again, we are interested in the specific case when $\vec{L} \approx \vec{V}$ (Figure 5.12.d). Under these conditions, the modified B-K model obtained from Equation (5.3) is

$$I_d(\theta_i) \approx \frac{\pi^3}{A}\left(\frac{4\sigma T}{\lambda}\right)^2 \cos^3\theta_i \exp\left\{-\left(\frac{2\pi}{\lambda}\right)^2 \left[(T\sin\theta_i)^2 + (2\sigma\cos\theta_i)^2\right]\right\} \qquad (5.28)$$

Here since λ appears in the exponential function, the normalized intensity does depend on λ while it does not depend on A. In Figure 5.11, we show the behavior of the modified B-K model given by Equation (5.28). The different curves in each panel show the normalized diffuse intensity as a function of the incidence angle for different values of surface roughness σ and correlation length T. In panel (a) we use a fixed value of T and several values of σ, whereas in panel (b) we use a fixed value of σ and several values of T. Here we use $\lambda = 700nm$. The dashed curve in each panel is the prediction of Lambert's

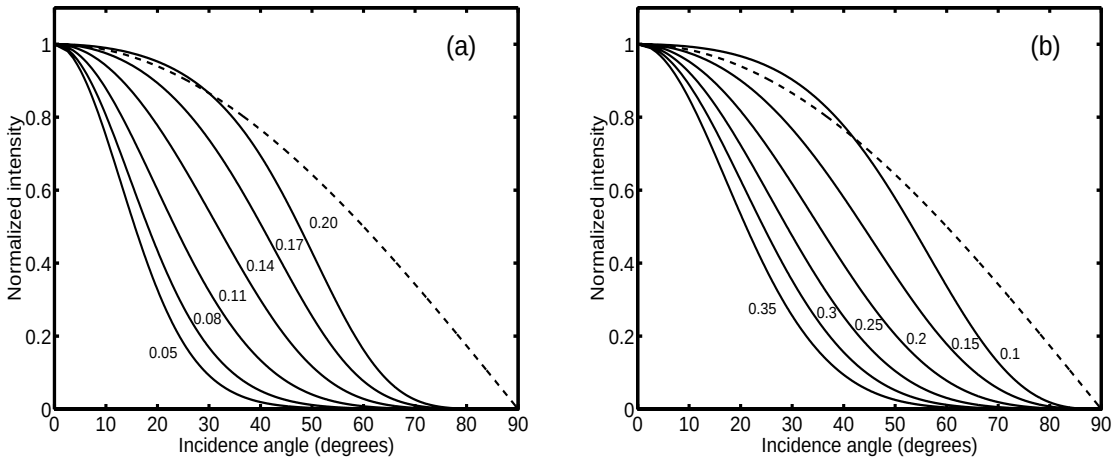

Figure 5.11: Normalized diffuse intensity using the modified B-K model (Equation 5.25) versus incidence angle compared to Lambertian model (dashed) for $\theta_i = \theta_r$ and $\phi_r = \pi$: (a) $T = 0.12\mu m$ and $\sigma = 0.05, ..., 0.2\mu m$, (b) $\sigma = 0.1\mu m$ and $T = 0.1, ..., 0.35\mu m$.

law. From the different curves it is clear that when the correlation length is fixed, the larger the surface roughness, the higher the diffuse intensity. In contrast, when the surface roughness is fixed, the larger the correlation length, the lower the diffuse intensity.

Lookup Table Solution

Neither of the simplified variants of the B-K model are analytically tractable and hence the Lambertian reflectance component can not be recovered directly in closed form. Instead, in a manner similar to that described earlier in Chapter 4, we adopt a lookup table approach as a practical alternative. To do this we tabulate $\cos\theta_i$ as a function of the computed diffuse intensity from Equations (5.27) and (5.28). Once the parameters of these models are estimated, then the lookup table may be computed using the corresponding equation. Under conditions where $\vec{L} \approx \vec{V}$ (Figure 5.12.d), the modified B-K model is amenable to our lookup table approach since, like Lambert's law, the brightness decreases monotonically with increasing incidence angle (Figures 5.10.a and 5.11). As a result the reflectance functions appearing in Equations (5.27) and (5.28) are injective and invertible. Note also, that for other illumination geometries where the reflectance is directly dependent on both the incidence and reflectance angles, the lookup table approach is not

necessarily applicable.

5.4.2 Surface Normal Recovery and Re-illumination

Once the corrected Lambertian image is to hand, then the surface normal directions may be recovered by applying a standard SFS method to the corrected image. Here we use our SFS method which is later described in Chapter 6. Once accurate surface normals are in hand, re-illuminating object surfaces is a straightforward task. Since at each point on the surface, the local surface normal and the viewing direction are known (and fixed), we can use the general variants of the modified B-K model to compute the diffuse intensity for any light-source direction.

5.5 Experiments

The experimental evaluation of the physics-based models for rough surface reflectance is divided into four parts. We commence in Section 5.5.1 by showing how the methods developed in Section 5.3 can be used to estimate the physical roughness parameters for a variety of different materials. With the surface parameters to hand, in Section 5.5.2 we show the results of fitting the B-K model and the Vernold-Harvey modification to intensity data for cylindrical objects. Section 5.5.3 takes the study one step further and shows how corrected Lambertian images may be recovered from rough surfaces. Finally, in Section 5.5.4 we show how SFS can be applied to the corrected Lambertian image to recover surface normal directions and demonstrate how these surface normal directions may be used for object re-illumination.

The experimental facilities at our disposal are relatively rudimentary. The images used in our experiments have been captured using an Olympus 10E camera. Each surface has been imaged under controlled lighting conditions in a darkroom. Different geometries of the experimental setups are shown in Figure 5.12. The objects have been illuminated using a single collimated tungsten light-source with a parallel beam whose mean wavelength is approximately $550nm$. The light-source was placed 2m away from the objects and

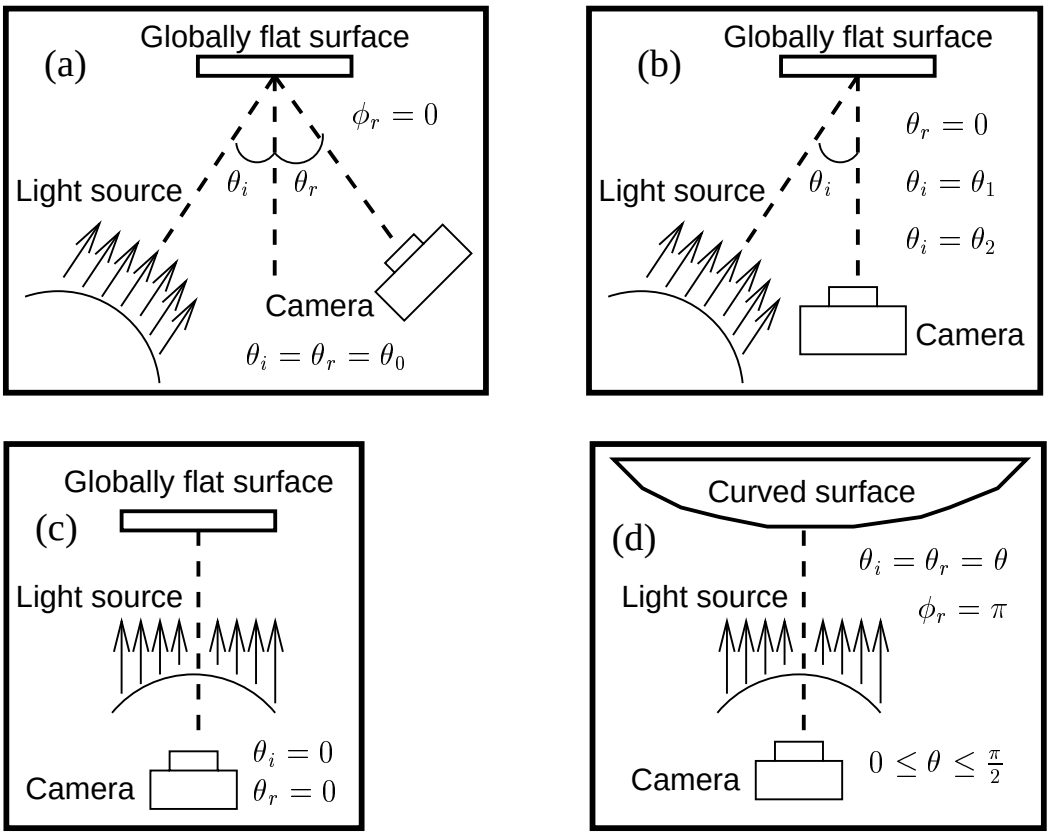

Figure 5.12: Experimental setup configurations: (a) measuring off-normal specular intensity to estimate σ for slightly-rough dielectrics; (b) measuring two diffuse intensity values to estimate T for slightly-rough or σ/T for very-rough dielectrics; (c) measuring normal specular reflectance to estimate σ for metals; (d) illumination conditions for photometric correction of curved surfaces.

collimated using a cylinder 0.5m long and with diameter 0.12m. The distance between the camera and the objects was 1.5m. The camera was focused automatically with a focal length between f/2.0 and f/2.4. We set the aperture ratio to f3.2. The camera (viewing) direction $\vec{V}$ is aligned with the z axis and the light-source direction $\vec{L}$ is recorded when the images are captured. Finally, to experiment under the conditions where $\vec{L} \approx \vec{V}$, the light-source and camera are slightly displaced in the vertical plane.

5.5.1 Rough Surface Parameter Estimation

Turning our attention first to metallic surfaces, we have performed measurements on flat bars composed of aluminium, iron and steel, and, also a pair of stainless steel scissors as examples of slightly-rough metallic surfaces. For each metallic surface we capture one

image under normal illumination. These images are shown in Figure 5.13 (panels a-d). Here we use both Equation (5.25) and Equation (5.26) to estimate the RMS roughness

Figure 5.13: Images of (a-d) metallic surface samples of aluminium, iron, steel and stainless steel scissors under normal incidence ($\theta_i = 0$); (e-g) dielectric surface samples of porcelain plate, plastic plate and glossy paper under off normal incidence ($\theta_i = \theta_r = 14^\circ$).

of these surfaces. The results are shown in Table 5.1. To evaluate the accuracy of these results, one should ideally compare them with ground-truth measurements of roughness based, for instance, on profilometry measurements. However, this is beyond our experimental resources. Hence, in Table 5.1, we show the typical values of RMS roughness for some of the metallic surfaces which have been tabulated in the literature (Depew and Weir, 1971). Obviously, these values are not for the samples we have used, but at least they show that our estimated values are of the correct order. The values of the roughness estimated using Equation (5.25) are in better agreement with the typical values than those obtained using Equation (5.26). The reason is that Equation (5.26) is based on the assumption that $\sigma \ll \lambda$ and the wavelengths used here do not satisfy this condition very well. Having estimated the RMS roughness, we can easily use the Torrance formulation (Equation 5.15) to estimate the RMS slope m. Here, the measured reflectance quantity R in Equation (5.15) contains both the diffuse and specular reflectance components as described in Section 5.3.3. The reflectance values are measured by averaging the image intensity over a neighborhood containing both the locations where the specular reflectance is dominant and neighboring locations where the diffuse reflectance is dominant (Figure 5.12.a). We also find the correlation length T using the formula $T = \sigma\sqrt{2}/m$. Again, we present the results in Table 5.1.

Next, we turn our attention to slightly-rough dielectric surfaces. Here we study a porcelain plate, a white plastic plate and glossy paper. For each of these surfaces we

Surface type	$\sigma(\mu m)$, Eq. (26)	$\sigma(\mu m)$, Eq. (25)	$\sigma(\mu m)$, typ.	m	$T(\mu m)$
Aluminium	0.0545	0.0988	0.100	0.0437	3.1943
Iron	0.0539	0.0927	N/A	0.0382	3.4374
Steel	0.0530	0.0852	0.086	0.0355	3.3958
Stainless steel	0.0551	0.1086	0.157	0.0526	2.9224

Table 5.1: Results for the slightly-rough metallic surfaces shown in Figure 5.13.

capture one image under off-normal illumination and in the specular direction (Figure 5.12.a). Here $\theta_i = \theta_r = 14^\circ$. These images are shown in Figure 5.13 (panels e-g). In this figure we have enhanced the contrast of the white plastic image by 80%, since the original image contained little intensity structure and was difficult to visualize. Here, we use the B-K model for the specular direction to find the surface roughness σ (Equation 5.18). We also estimate the correlation length T using the technique outlined in Section 5.3.1 (Equation 5.20). In doing this, we capture two images illuminated from different directions, i.e. I_1 at $\theta_i = 30^\circ$ and I_2 at $\theta_i = 45^\circ$ (Figure 5.12.b). We have not shown these images here since there is little visible difference between them. The main reason for this is that for these surfaces the irregularities and topographic features are not visible by naked eye. We present the resulting estimates of RMS roughness σ, correlation length T and also their ratio σ/T in Table 5.2.

Surface type	i_{spec}	σ (μm)	I_1/I_2	T (μm)	σ/T
Glossy paper	0.013972	0.119	1.04270	0.246	0.4837
Porcelain plate	0.003324	0.137	1.04716	0.323	0.4241
Plastic plate	0.001702	0.145	1.03179	0.349	0.4155

Table 5.2: Results for the slightly-rough dielectric surfaces shown in Figure 5.11.

Finally, we present results for very-rough dielectric surfaces. The samples used in our study are stone tiles, textured wall papers, sandpapers and a wooden cube. For each of these surfaces, we capture two images under two different incidence angles (Figure 5.12.b), i.e. I_1 at $\theta_i = 30^\circ$ and I_2 at $\theta_i = 45^\circ$. These image pairs are shown in Figure 5.14. We also show the difference between the two images for each surface at the bottom row of the figure. We have enhanced the contrast of these difference images by 50%. Here

the darker the point in a difference image, the higher the difference.

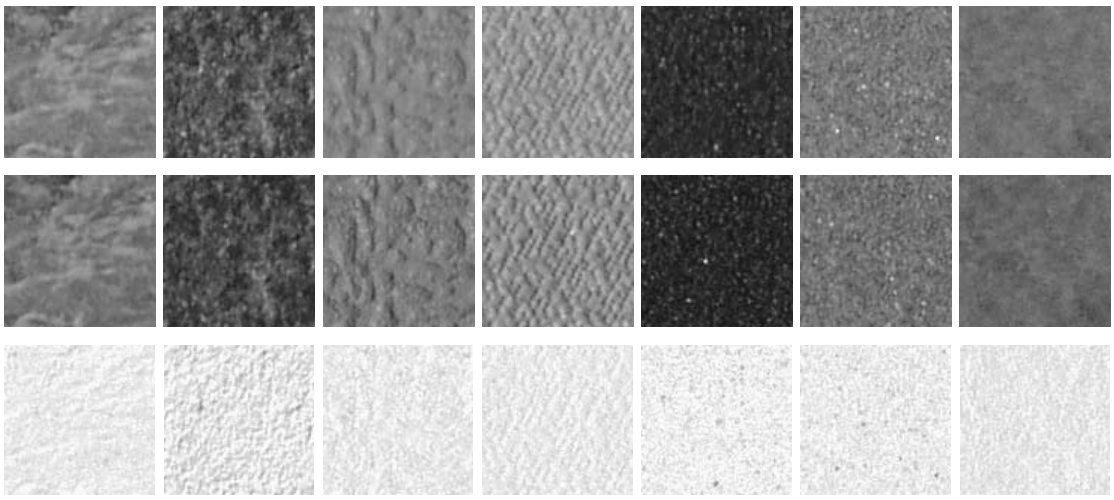

Figure 5.14: Samples of (a-b) stone tiles, (c-d) textured wall papers, (e-f) sandpapers, and (g) wooden cube for (1) $\theta_i = 30^\circ$ and (2) $\theta_i = 45^\circ$; (3) enhanced difference images.

For these surfaces we use the method explained in Section 5.3.2 using both variants of the modified B-K model to find the corresponding ratios σ/T_G and σ/T_E (Equations 5.22 and 5.24). The results are presented in Table 5.3. For surfaces with an identical assumed height distribution function (Gaussian) and two different correlation functions, i.e. the Gaussian and the exponential, we expect to find two different values of σ/T for each surface. The reason is that while σ is unaffected by the changes of correlation function, the different correlation functions will result in different estimates for the correlation length, i.e. T_G and T_E. In order to gauge the validity of the model, for very-rough surfaces with a Gaussian correlation function we can also compute the angle of RMS slope θ_K, which gives an upper limit for the valid incidence angles, and the energy conservation rate ξ_K, which measures the fraction of the energy that is not absorbed by the surface.

There are several points to note concerning the results given in Tables 5.1-5.3. First, the magnitude of the slope parameter σ/T can be used to distinguish between slightly-rough metals and dielectrics, and, very-rough dielectrics. Also, while the RMS roughness σ for slightly-rough metals and dielectrics are comparable, the correlation lengths T differ by an order of magnitude. Hence, the measured parameters could in principle be used to

Surface type	I_1/I_2	σ/T_G	$\theta_K(^\circ)$	$\xi_K(\%)$	σ/T_E
Stone tile 1	1.056	0.9126	38	25.9	2.2080
Stone tile 2	1.141	0.4833	56	65.7	1.1315
Wall paper 1	1.077	0.7104	45	39.1	1.7050
Wall paper 2	1.070	0.7654	43	34.7	1.8420
Sandpaper 1	1.091	0.6332	48	46.4	1.5113
Sandpaper 2	1.088	0.6483	47	44.8	1.5494
Wooden cube	1.079	0.6964	45	40.3	1.6697

Table 5.3: Results for the very-rough dielectric surfaces shown in Figure 5.14.

classify the surface type. However, this topic is outside the scope of our study.

5.5.2 Comparing the Beckmann model with Reflectance Data

With the roughness parameters to hand, we can compare the predictions of the B-K model with measured intensity data for curved objects. In order to simplify the comparison, we work with cylindrical objects since we can estimate the incidence angles using simple geometry as described in Chapter 3 (Section 3.7.6). Hence, we have constructed cylinders using the two samples of sandpaper (types 1 and 2) previously used to estimate the values of σ/T (see Table 5.3). We illuminate each cylinder so that $\vec{L} \approx \vec{V}$ (Figure 5.12.d). The B-K model applicable to this case is summarized by Equations (5.6) to (5.9). We use the estimated values of σ/T for the samples of sandpaper to compare the predictions of the model with the data for the cylinders in the following way. Across a horizontal line, perpendicular to the axis of each cylinder, we compute the mean-intensity at each point (on the right-hand half-cylinders) by averaging the intensity values over the vertical neighbors. The mean-intensities computed in this way are shown as a function of the incidence angles (computed from the known cylindrical geometry) in Figure 5.15 as solid curves. In this figure, panels (a) and (b) are for the sandpapers of type 1 and 2, respectively. Here, we show the Lambertian model (dotted curve), the original B-K model with a Gaussian correlation function (Equation 5.6, dash-dot curve), and the Vernold-Harvey modified B-K model with both a Gaussian (Equation 5.8, dash-dot curve) and an exponential (Equation 5.9, dashed curve) correlation function. In both plots, only for $\theta_i < 10^\circ$, the data are close

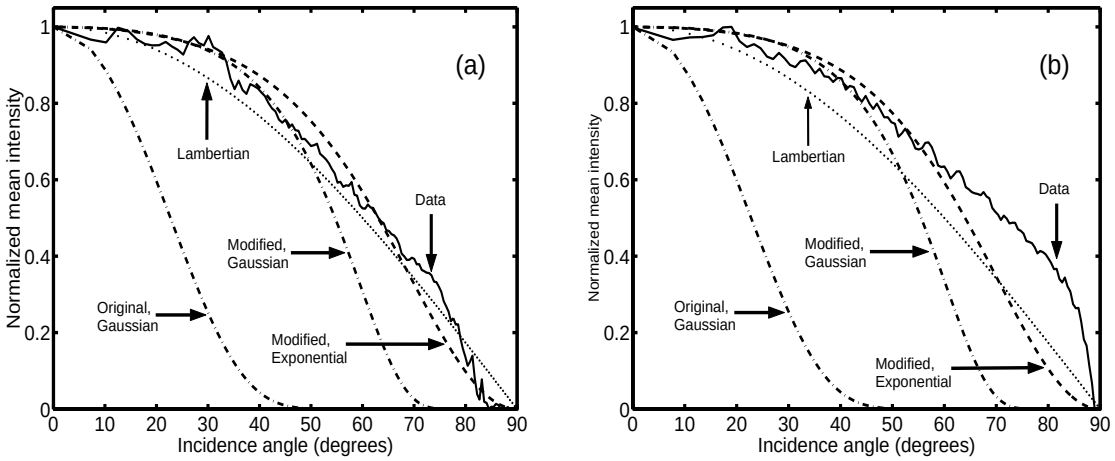

Figure 5.15: Normalized intensity (solid) on sandpaper cylinders (a) type 1 and (b) type 2 versus incidence angle compared to the variants of the B-K model and Lambert's law.

to the predictions of the original B-K model. However, and as expected, they deviate significantly from this model as the incidence angle increases. In panel (b), for incidence angles less than 60° the data curve fits well to both variants of the modified B-K model, while it deviates significantly from the Lambertian model. In panel (a), on the other hand, there is good agreement between the data and both variants of the modified B-K model at most incidence angles. However, in both plots, for larger incidence angles, the modified B-K model with an exponential correlation function is much closer to the data than the remaining models. This supports the claims in the literature (Bennett and Mattsson, 1999; Ogilvy, 1991) that the exponential correlation function gives a better fit to measured surface data. From these plots, it is also clear that the Vernold-Harvey modification results in good agreement between the modified B-K model and the real-world data.

As noted earlier the failure of the models at large angles is attributable to multiple scattering effects (O'Donnell and Mendez, 1987) and, as noted by Caron et al. (Caron et al., 2002), can be characterized using the energy conservation rate, ξ_K, and the upper limit for the incidence angles, θ_K. For the two samples of sandpaper used here, the energy conservation rates $\xi_K = 46.40\%$ and $\xi_K = 44.83\%$ (Table 5.3) suggest the possibility that the scattered energy may have been misestimated by as much as 50% when the Kirchhoff theory is used. Moreover, $\theta_K = 48.16^\circ$ and $\theta_K = 47.48^\circ$ (Table 5.3) are very close to the

incidence angles given by Figure 5.15 above which the Kirchhoff theory is expected to fail. Specifically, the angle at which the curve corresponding to the modified B-K model (with a Gaussian correlation function) moves away from the data curve is about 46° in Figure 5.15.a and 44° in Figure 5.15.b.

To conclude from these experiments, we found that for dielectric surfaces with different scales of roughness, it is the modified B-K model which results in reasonable predictions of reflectance for most angles of incidence and reflectance. Also, for the very-rough dielectric surfaces, the modified B-K model with an exponential correlation function gives the best agreement with the real-world data. Note also that, we have not tested the predictions of the Davies theory against real-world data.

5.5.3 Photometric Correction

In Section 5.5.2 we established that the Vernold-Harvey modification of the B-K model gives reasonable agreement with measured intensity data over a significant range of incidence angles. We now take our experimental study one step further by showing how this modification can be used to extract corrected Lambertian reflectance from rough surfaces.

Very-Rough Surfaces

The objects used in our study are composed of terra-cotta and sandpaper (a bear and cylinder) and are shown in Figure 5.16. To perform photometric correction, we must first estimate surface slope $m = \sigma/T$. We use the technique outlined in Section 5.3.2. The estimated values are $m = 1.41$ for the sandpaper and $m = 1.29$ for the terra-cotta. Using these parameters we construct a lookup table as detailed in Section 5.4.1, using Equation (5.27). We present the results in Figure 5.16. Here, panel (a) shows the raw images I, panel (b) shows the recovered Lambertian images I_L and panel (c) shows the difference images ($|I - I_L|$). Here the darker the point in a difference image, the larger the difference. We use the absolute function to show both positive and negative differences. It can be seen in Figure 5.10.a that for $m = 1.3$ (terra-cotta) if $\theta_i < 45^\circ$ then $I_L < I$,

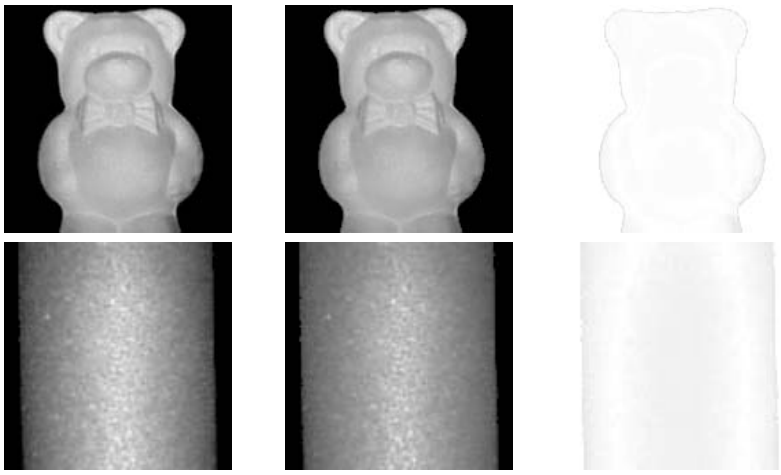

Figure 5.16: Photometric correction on images of (1) a terra-cotta bear and (2) a sandpaper cylinder: (a) raw images; (b) corrected Lambertian images; (c) difference images.

whereas for $m = 1.4$ (sandpaper), if $\theta_i < 60^\circ$ then $I_L < I$. Although the correction process has an effect at almost every location on each surface, the differences are most marked where the inclination of the object surface is steepest.

In Figure 5.17 we display some scatter plots to emphasize the effects of the photometric correction process for the images in Figure 5.16. In panel (a) we plot the Lambertian intensity I_L versus the original (raw) intensity I. In panel (b) we plot the quantity $(I_L - I)$ versus I. The plots in the top row are for the terra-cotta bear, whereas the plots in the bottom row are for the sandpaper cylinder. In the case of the terra-cotta bear, it is clear that the corrected Lambertian intensity is almost always greater than the raw image intensity. However, in the case of the sandpaper cylinder, there are intensity intervals over which the corrected Lambertian intensity is both smaller and larger than the raw intensity.

Slightly-Rough Surfaces

The slightly-rough objects, the bear and the hand, are made of porcelain and are shown in Figure 5.18. The estimates of the RMS roughness σ and the correlation length T required to apply Lambertian correction to these objects are taken from Table 5.2 (porcelain plate). We construct a lookup table using the model of Equation (5.28) where $\sigma = 0.137\mu m$ and $T = 0.323\mu m$. For slightly-rough surfaces both specular and diffuse reflectance

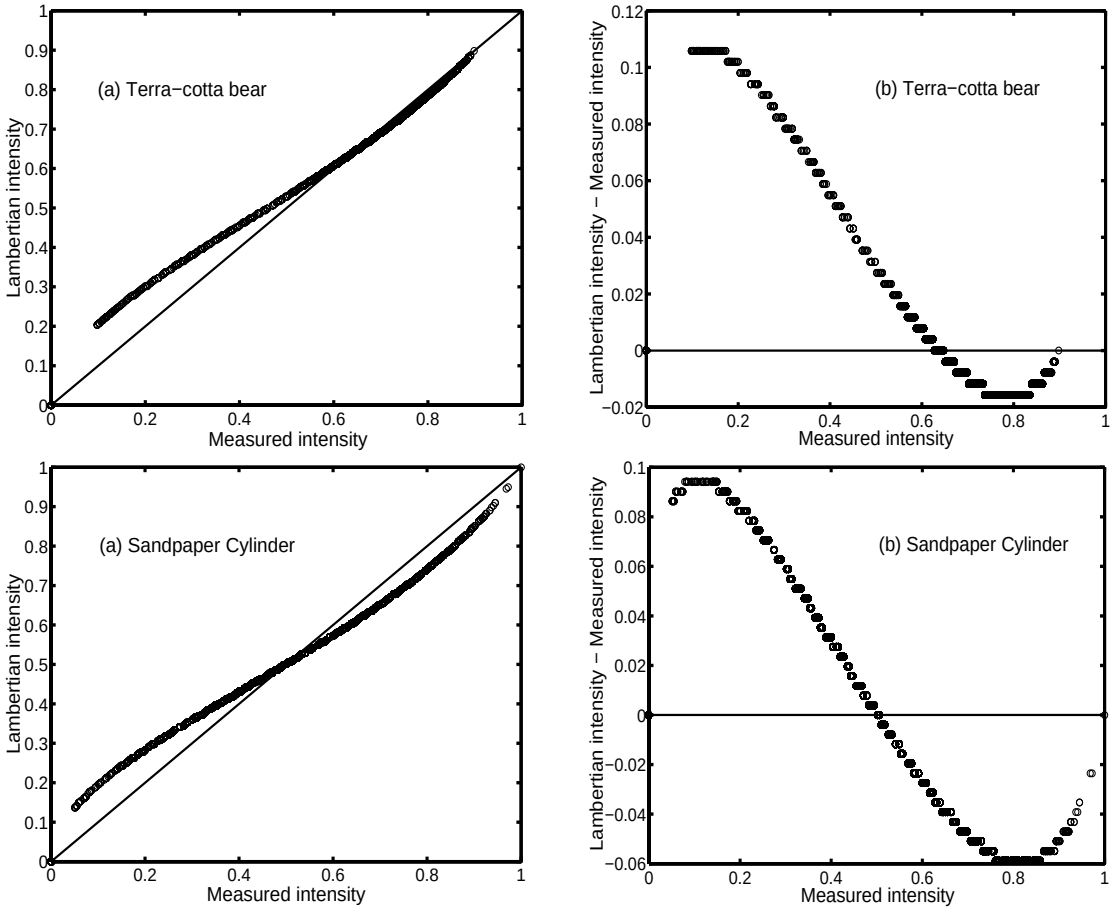

Figure 5.17: Photometric correction for (1) the terra-cotta bear and (2) the sandpaper cylinder: (a) Lambertian intensity I_L versus measured intensity I; (b) $(I_L - I)$ versus I.

components are present. The specular component must therefore be removed before we perform our photometric correction on the diffuse reflectance component. This is done using the method that we have described in Chapter 3. In Figure 5.18, panel (a) shows the raw image. In panel (b), we show the specular intensity reconstructed, using the Torrance-Sparrow model (Torrance and Sparrow, 1967). Panel (c) shows the result of performing specularity removal as described in Chapter 3. These are the diffuse images I_D used as inputs to the correction process. Panel (d) shows the recovered Lambertian images I_L and those in panel (e) show the difference images $(I_L - I_D)$. Here the brighter the point in a difference image, the larger the difference.

In Figure 5.19, in a manner similar to Figure 5.17, we show the scatter plots for the photometric correction results on the porcelain bear and the porcelain hand. Here, at most

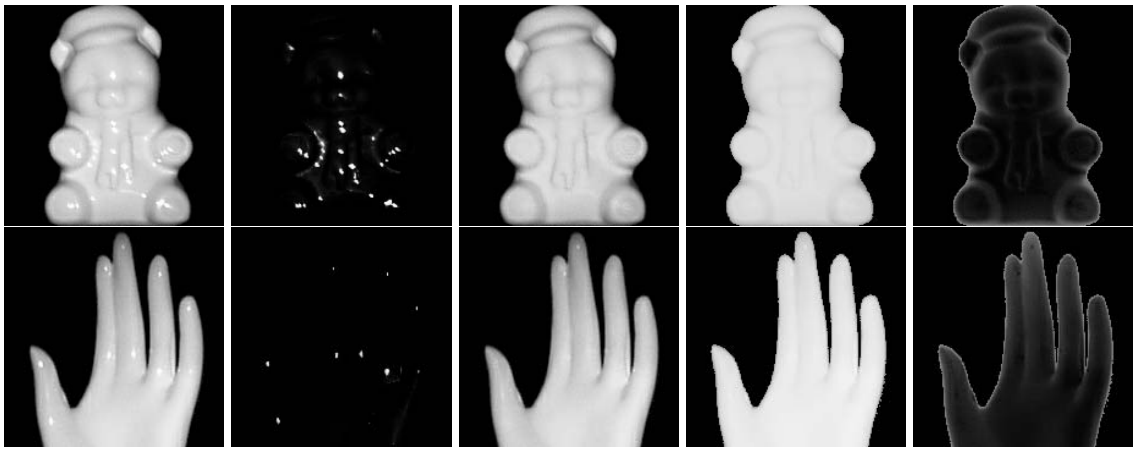

Figure 5.18: Removing (b) the Specular components from (a) raw images to obtain (c) diffuse images using the method of Chapter 3; (d) using the modified B-K model to recover Lambertian images; (e) differences between the Lambertian and diffuse images.

image locations where specular reflectance is absent, the recovered Lambertian intensity shows a reasonable increase over the original intensity, and, the increase is largest for low intensities. It is clear from the original images in Figure 5.18 that the points with the largest brightness correspond to the specular highlights. The Lambertian intensities for these specific points are smaller than the corresponding original intensities. The reason for this is that at the locations of specular highlights, the corrected grey-level values are the results of two different correction methods. Firstly, a highlight removal process is applied to the raw image, and secondly a lookup table correction method is applied to the separated matte image. In comparison with the plots in Figure 5.17, here the increase in image intensity is on average larger. Also, most of the locations on both porcelain objects are affected by an intensity increase rather than an intensity decrease.

5.5.4 Shape-from-Shading and Re-illumination

The corrected Lambertian images obtained in the previous section may be used as input images to the SFS method which is later described in Chapter 6. The needle maps of the surface normals obtained using this method are shown in Figures 5.20-5.21. In each figure, panel (a) shows the surface normals extracted from diffuse images whereas panel

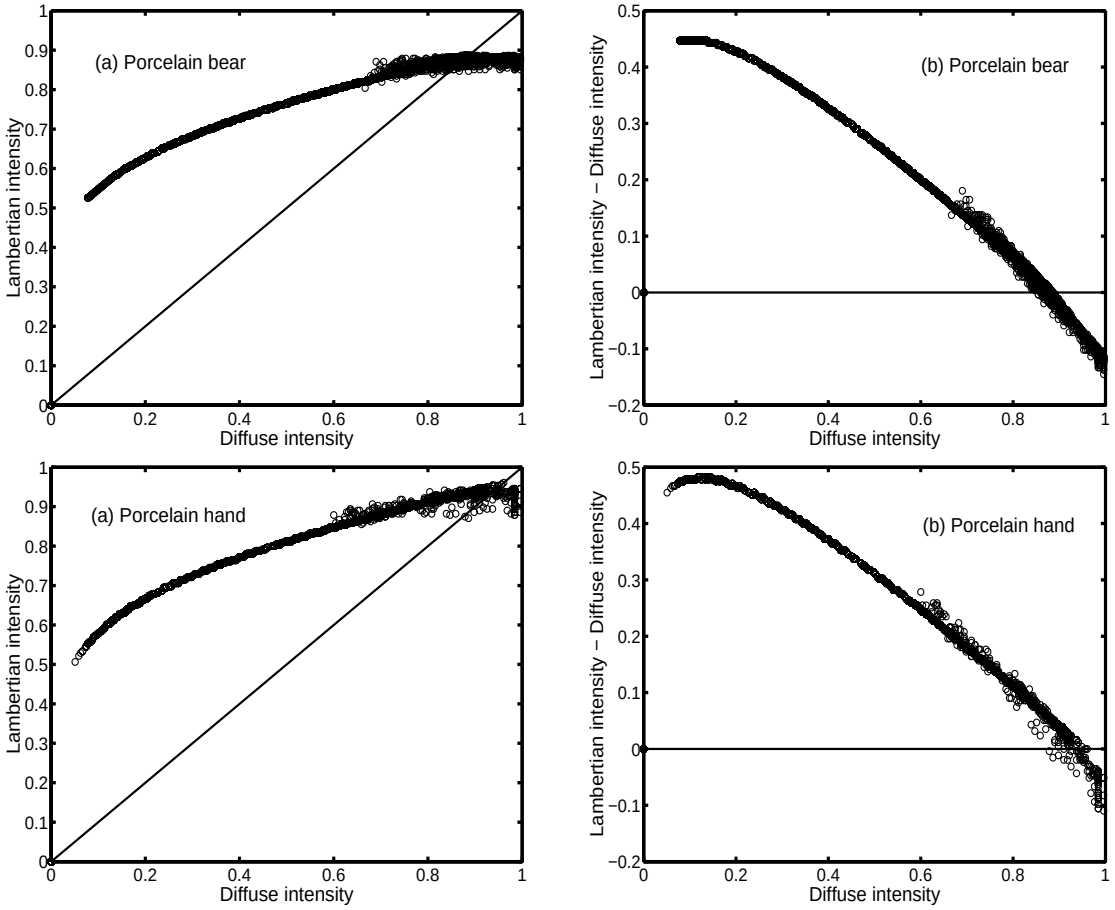

Figure 5.19: Photometric correction for (1) the porcelain bear and (2) the porcelain hand: (a) Lambertian intensity I_L versus diffuse intensity I; (b) $(I_L - I)$ versus I.

(b) shows those extracted from Lambertian images. The field of difference between these surface normals is shown in panel (c). Figure 5.20 is for the very-rough surfaces (Figure 5.16) , while Figure 5.21 is for the slightly-rough surfaces (Figure 5.18). In Figures 5.20-5.21, the complex surface structure of the terra-cotta bear and the porcelain bear are clearly visible. For the sandpaper and the porcelain fingers, the symmetrical shapes are nicely preserved. Note also that since the specular highlights are removed from the raw images of porcelain objects, no corresponding artifacts appear on the needle maps.

We now turn our attention in more detail to the effects of the Lambertian correction process on the surface normal directions. The surface normals are constrained to fall on a cone whose opening angle is given by the inverse cosine of the image intensity. The Lambertian correction causes changes both in the zenith angles and in the azimuth angles

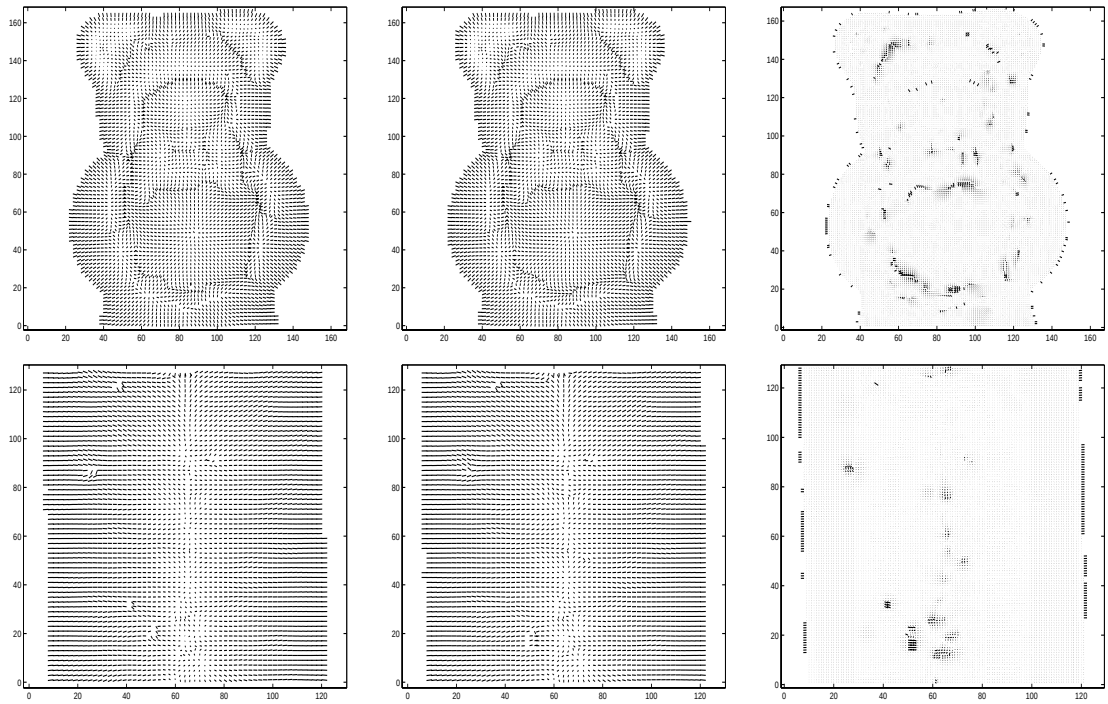

Figure 5.20: Surface normals obtained by applying Lambertian SFS (Chapter 6) to the (a) diffuse and (b) Lambertian images of Figure 5.16; (c) field of differences in needle maps.

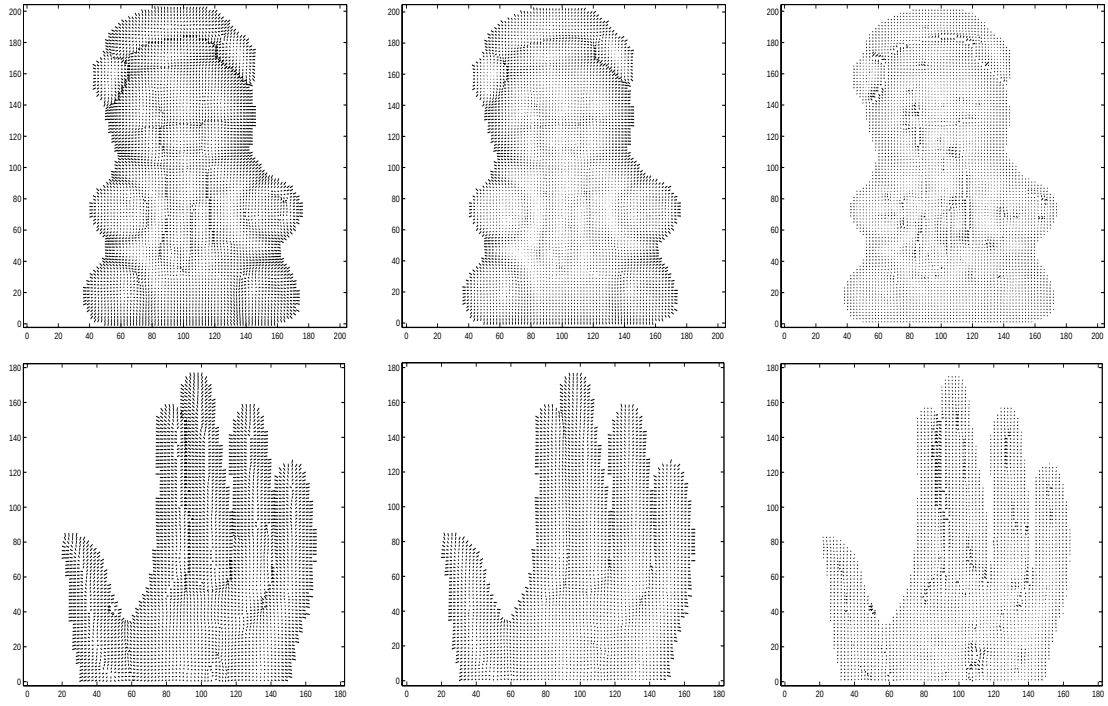

Figure 5.21: Results (as described in Figure 5.20) for the images of Figure 5.18.

of the surface normals. In qualitative terms, when the image becomes brighter then the apex (opening) angle of the Lambertian cone becomes smaller. Hence the zenith angle of the surface normal decreases. When the image becomes darker, on the other hand, then the zenith angle increases. To illustrate the effect of the Lambertian correction process, we have produced scatter plots for the azimuth and zenith angles for surface normals extracted from the diffuse images I and the corresponding corrected Lambertian images I_L. In Figure 5.22, the top row is for the terra-cotta bear whereas the bottom row is for the sandpaper cylinder. Here, panel (a) shows the plot for the zenith angle, i.e. $\cos^{-1}(I_L)$

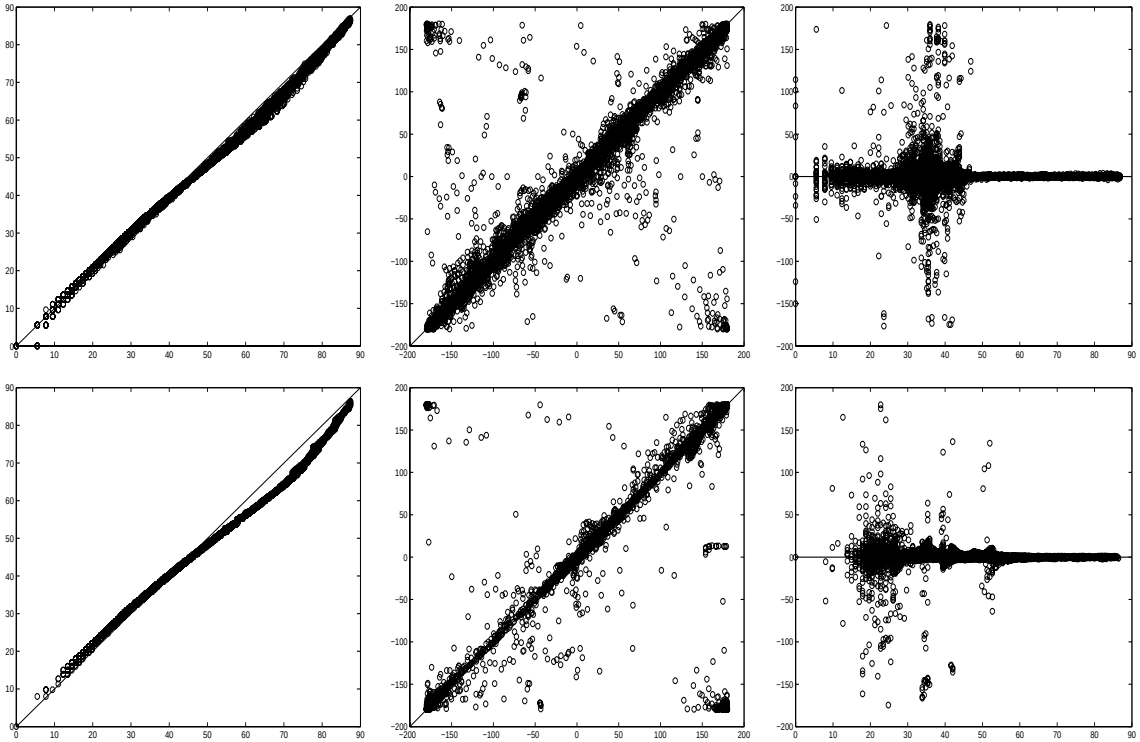

Figure 5.22: Plots for the incidence (a) zenith and (b) azimuth angles (degrees) extracted from the recovered surface normals for (1) the terra-cotta bear and (2) the sandpaper cylinder: the horizontal axis shows the angles corresponding to the diffuse image, while the vertical axis shows those corresponding to the Lambertian image; (c) the difference between the diffuse and Lambertian azimuth angles versus the Lambertian zenith angle.

versus $\cos^{-1}(I)$, while panel (b) shows the plot of the azimuth angle. In both plots, the horizontal axis shows the angle extracted from the surface normals corresponding to I while the vertical axis shows the angle extracted from those corresponding to I_L. In the zenith angle plot, there is no scatter and the data points trace out a curve. The reason for

this is that the Lambertian SFS method constrains the surface normals to fall on a cone whose apex angle is the arc-cosine of the image intensity. The axis of the cone points in the light-source direction, and since $\vec{L} \approx \vec{V}$, the apex angle of the cone is equal to the zenith angle. The azimuth angle plot, however, does exhibit scatter. The reason for this is that the surface normals are free to rotate about the cones subject to smoothness constraints. However, in the azimuth plot there is a clear regression line.

To investigate the stability of the surface normals under changes in azimuth angle, panel (c) of Figure 5.22 shows the difference in azimuth angle for the surface normals computed using I and I_L as function of the Lambertian zenith angle $\cos^{-1}(I_L)$. It is clear that the largest differences in azimuth angle occur at intermediate zenith angles. Hence, these are associated with moderate changes in surface normal direction. This means that the modified B-K model results in significant differences in surface normal directions.

Finally, using the surface normals extracted from the recovered Lambertian images, we experiment with re-illuminating objects from different directions. In Figure 5.23, we show re-illuminations of the surface normals for the very-rough surfaces (Figure 5.20), and in Figure 5.24 for the slightly-rough surfaces (Figure 5.21). In Figure 5.23, we compute the total (or diffuse) reflectance using the modified B-K model given by Equation (5.5), where F^2 is replaced by $\cos\theta_i$. Here, in both rows from panel (a) to (e), the light-source directions are $(18^\circ, -90^\circ)$, $(9^\circ, -90^\circ)$, $(0, 0)$, $(9^\circ, 90^\circ)$ and $(18^\circ, 90^\circ)$ respectively, where the first and second numbers are the values of the zenith and azimuth angles. In Figure 5.24, we compute the diffuse reflectance using the modified B-K model appearing in Equation (5.3), where again F^2 is replace by $\cos\theta_i$. The light-source directions used here are different from those used in Figure 5.23 in order to show the reconstructed images for a wider range of re-illumination angles. Here, the light-source directions from panel (a) to (e) are $(40^\circ, -90^\circ)$, $(20^\circ, -90^\circ)$, $(0, 0)$, $(20^\circ, 90^\circ)$ and $(40^\circ, 90^\circ)$, respectively. The images reflect the underlying shape of the objects in a consistent manner. However, for the reconstructed images of the porcelain hand, there is a "pimple" on the palm. This occurs at the location of a saddle on the surface, and is due a failure of the SFS method to resolve the concave/convex ambiguity where the principal curvature direction is indeterminate.

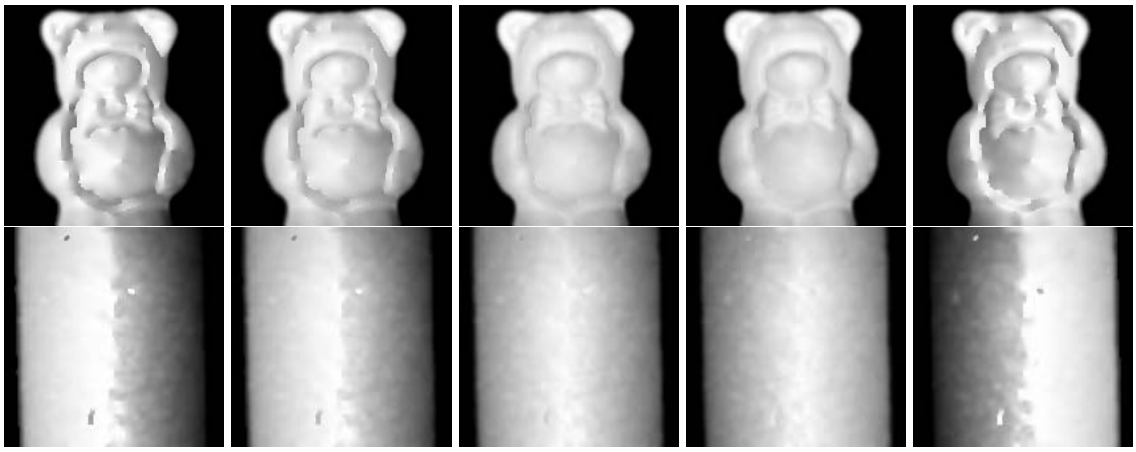

Figure 5.23: Re-illumination of the recovered surface normals (Figure 5.20) to generate new images using the modified B-K model for very-rough surfaces for the light-source directions (a-e): $(18°, -90°)$, $(9°, -90°)$, $(0, 0)$, $(9°, 90°)$ and $(18°, 90°)$.

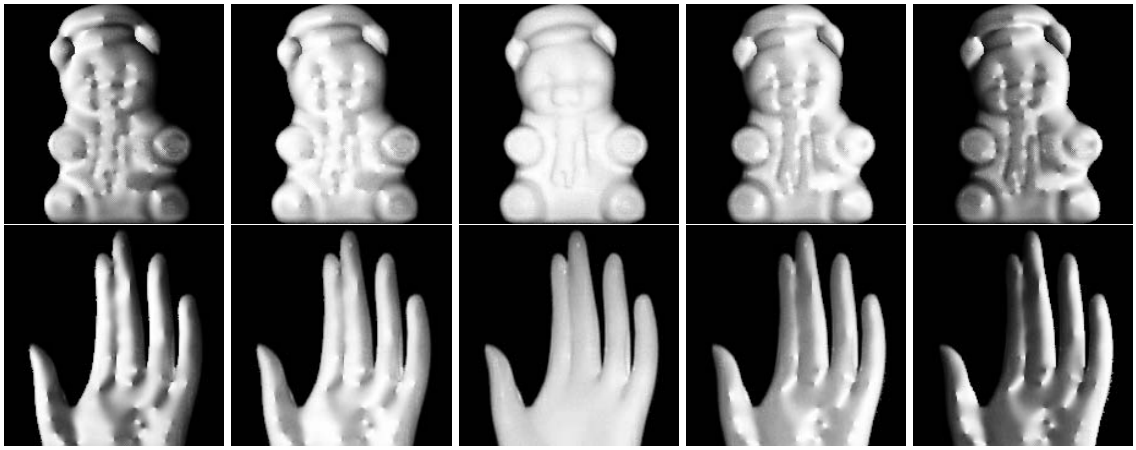

Figure 5.24: Re-illumination of the recovered surface normals (Figure 5.21) to generate diffuse images using the modified B-K model for slightly-rough surfaces for the light-source directions (a-e): $(40°, -90°)$, $(20°, -90°)$, $(0, 0)$, $(20°, 90°)$ and $(40°, 90°)$.

5.5.5 Physics-Based Rendering

Finally, we use a synthetic surface of a vase to show the results of rendering different images for a synthetic surface illuminated under different light-source directions using the modified B-K model for very-rough or for slightly-rough surfaces. In Figure 5.25 we show the ground-truth for the synthetic vase. In panel (a) we show the height information, whereas in panel (b) the field of surface normal directions is shown. The surface is a

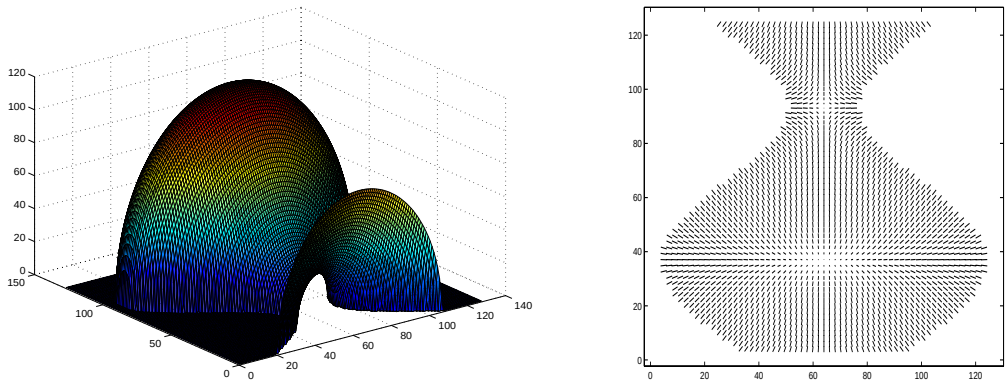

Figure 5.25: The ground-truth data of a synthetic vase which is a cylinder with a sinusoid section: (a) the surface height; (b) the surface normal directions.

cylinder with a sinusoid section.

Very-Rough Surfaces

In Figure 5.26, the different columns show diffuse images when the light-source direction is set to a fixed direction (α, β). Here α is the zenith angle and β is the azimuth angle of the light-source direction, while the viewing direction is set in the direction of z axis. In Figure 5.26, from the left-hand column to the right-hand column, we set (α, β) to $(0, 0)$, $(20°, 90°)$, $(50°, 90°)$ and $(90°, 90°)$. The top row of Figure 5.26 shows the images rendered using the Lambertian model. The remaining rows show the images obtained by rendering the surface using the modified B-K model (Equation 5.5, where F^2 is replaced by $\cos\theta_i$) with different surface slopes (m). In the second and third rows of the figure the surface slopes are set to $m = 1.0$ and $m = 1.2$. In panel (a) of these two rows, the light-source and viewing directions are identical. For these two cases, the modified B-K model results in lower reflectance compared to the Lambertian reflectance at every image location (Figure 5.10.a). However, in panel (d) where the light-source direction is perpendicular to the viewing direction, the pattern is not so straightforward. Here, for locations on the surface where the surface normal is in the same plane as the light-source and viewing directions (so that $\theta_i = \pi/2 - \theta_r$ and $\phi_r = 0$), the pattern may be reversed (Figure 5.10.b). For instance, the neck of the vase which has a cylindrical shape is the

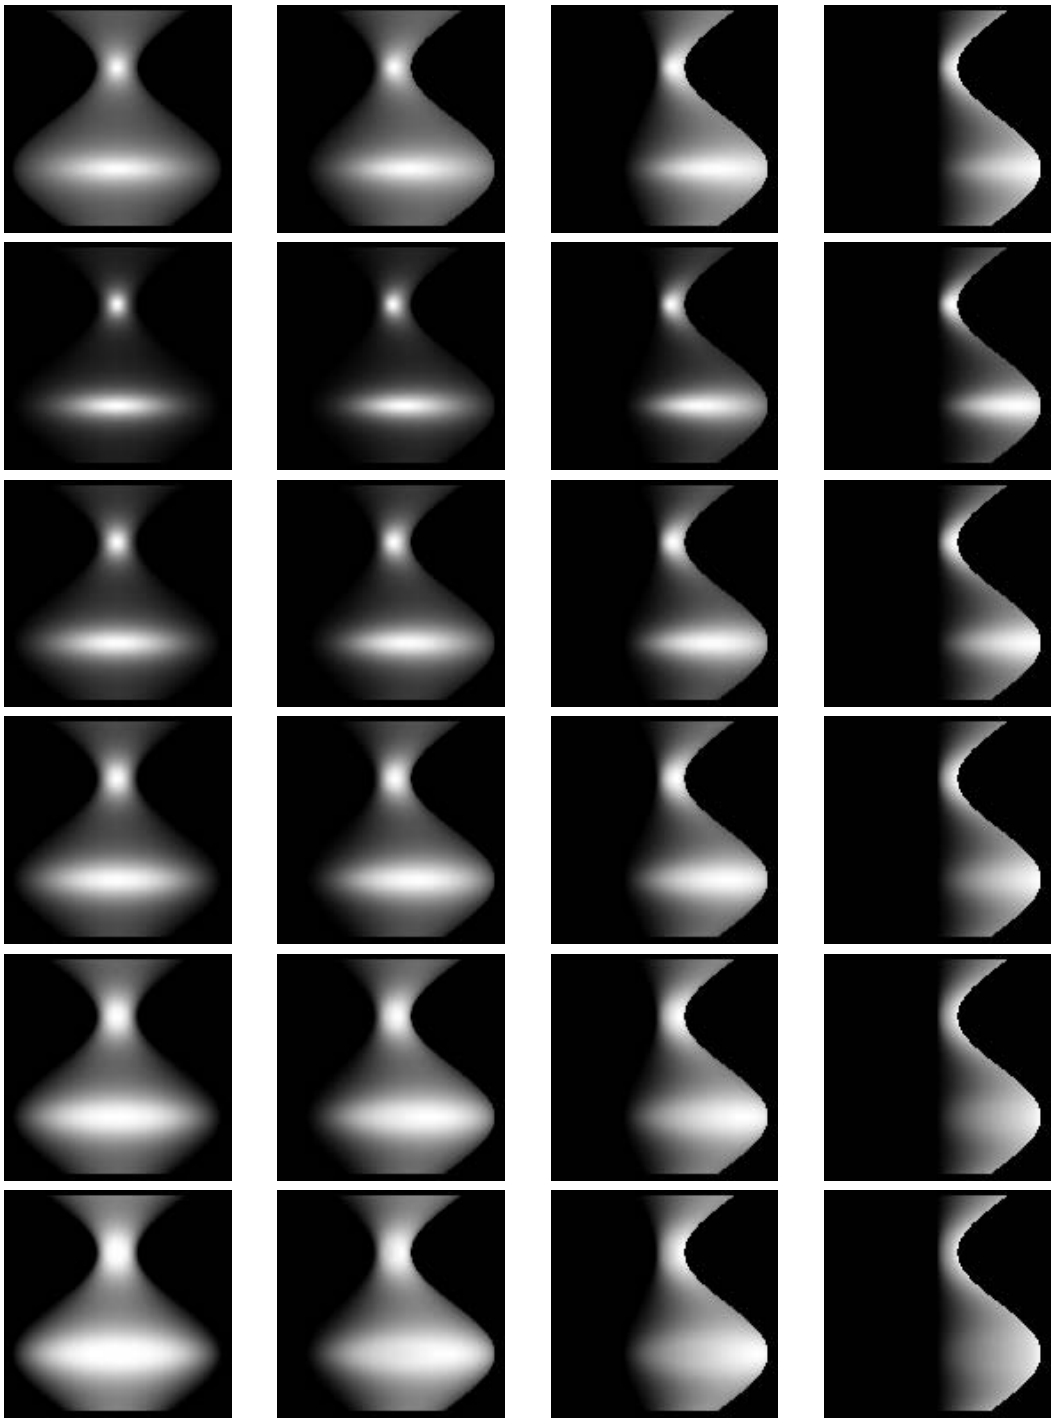

Figure 5.26: Surface rendering using (1) the Lambertian model and (2-6) the modified B-K model for very-rough surfaces with an exponential correlation function for the light-source zenith angles (a-d) 0°, 20°, 50°, and 90°; surface slope increases from the second row to the bottom row ($m = 1.0, 1.2, 1.4, 1.6, 1.8$).

main location for such surface normals. On the bulb of the vase where the radius of each cylindrical section is determined using a sinusoidal function, the only such locations are

at the section with the maximum radius. In the third, fourth and fifth rows, the surface slopes $m = 1.4$, $m = 1.6$ and $m = 1.8$ are used. For images in panel (a), as the surface slope value increases the portion of the surface for which the diffuse reflectance exceeds the Lambertian reflectance increases. Specifically, for $m = 1.8$ and for incidence angles less than $80°$, the diffuse reflectance exceeds the Lambertian reflectance (Figure 5.10.a). However, for the images in panel (d), for every point on the neck and specific points on the bulb, (i.e. the cylindrical section with the maximum radius) the pattern is reversed, i.e. the Lambertian reflectance exceeds the diffuse reflectance (Figure 5.10.b). For the remaining images which are shown in panels (b) and (c), the reflectance patterns can only be interpreted using the modified B-K model by considering the geometry of the tangent plane and the illumination direction.

Slightly-Rough Surfaces

Using the ground-truth surface normals of the synthetic vase, we experiment with rendering surface object from five different directions. At each point on the surface, we compute the diffuse reflectance using the modified B-K model for slightly-rough surfaces (Equation 5.3, where F^2 is replaced by $\cos\theta_i$). In Figure 5.27, we show rendered images when the light-source direction is re-positioned. The top row of Figure 5.27 shows the images rendered using the Lambertian model. The remaining rows show the images obtained by rendering the surface using the modified B-K model with different correlation length values. The RMS roughness is set to $\sigma = 0.1\mu m$ while the correlation length is increased from the second row to the bottom row, i.e. $T = 0.1, 0.15, 0.20, 0.25, 0.30\mu m$. From the left column to the right column, the light-source directions are $(0, 0)$, $(30°, 90°)$, $(60°, 90°)$ and $(90°, 90°)$ respectively, where the first and second numbers are the values of the zenith and azimuth angles. The images reflect the underlying shape of the object in a consistent manner. In a manner similar to that used earlier for Figure 5.26, the images of Figure 5.27 may also be interpreted using corresponding curves in Figure 5.11.b.

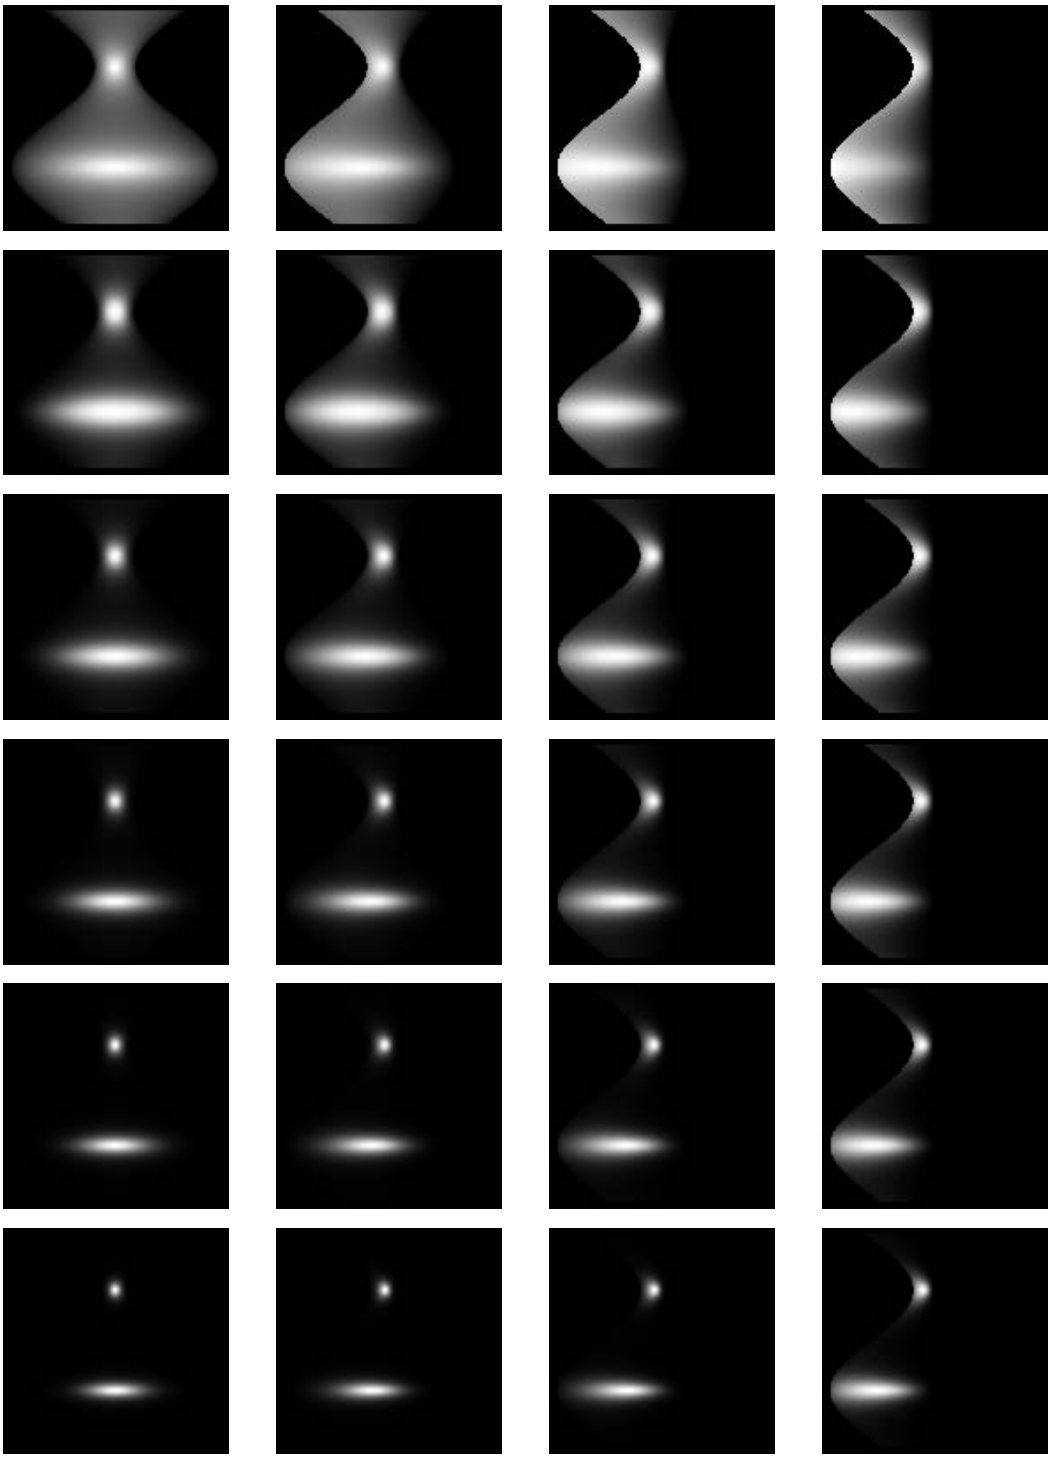

Figure 5.27: Surface rendering using (1) the Lambertian model and (2-6) the modified B-K model for slightly-rough surfaces for the light-source zenith angles (a-d) $0°$, $30°$, $60°$, and $90°$; the RMS roughness is fixed ($\sigma = 0.1\mu m$) while the correlation length increases from the second row to the bottom row ($T = 0.1, 0.15, 0.20, 0.25, 0.30\mu m$).

5.6 Conclusions

The main contribution in this chapter has been to explore the use of a family of reflectance models from physical optics for rough surface analysis tasks in vision. We have investigated two classes of model. The first of these is the Beckmann-Kirchhoff theory and its variants due to Vernold and Harvey (Vernold and Harvey, 1998), and, Caron et al. (Caron et al., 2002). The second is the Davies scattering theory, which is applicable to rough conducting surfaces. We have provided a review of the reflectance models and have commented on their domain of applicability. With the models to hand, we have considered how they may be applied to a number of different surface analysis tasks.

The first problem addressed was that of developing simple methods for estimating the physical parameters of rough surfaces using just a light-source and a digital camera. For slightly-rough surfaces, the estimated RMS roughness values are in good agreement with the typical values reported in literature. For very-rough surfaces, we have estimated the ratio of the RMS roughness to the correlation length. This estimated ratio allows the B-K model to be fitted to real-world data. While the modified B-K model of Vernold and Harvey fits well to the data, the original B-K model only fits well to the data for small incidence angles. The results are encouraging, since they indicate that physics-based reflectance models can be used in conjunction with simple experimental instruments, i.e. the visible light and a digital camera, to make surface roughness measurements.

Using the estimated values of the rough surface parameters, the second aim has been to explore whether the B-K model can be used to analyze surface shape. To embark on this endeavor, we have shown how the B-K model can be used to recover corrected Lambertian reflectance from curved surfaces. This allows us to apply standard SFS methods to rough surfaces and hence recover surface normal directions and also perform novel view synthesis through re-illumination.

Chapter 6

Curvature Consistency using Darboux Smoothing

In previous chapters of the thesis we focused on how reflectance models can be applied to surfaces to extract corrected Lambertian reflectance. The reason for doing this is to allow standard Lambertian shape-from-shading (SFS) methods to be applied to a greater variety of real-world surfaces. Specifically, in Chapter 3, we outlined a specular SFS algorithm for shiny surfaces so that the specular and diffuse reflectance components can be separated. In both Chapter 4 and Chapter 5, we utilized phenomenological and physics-based reflectance models to recover Lambertian component from the diffuse component.

In this chapter we describe our principal modification of the SFS method of Worthington and Hancock (Worthington and Hancock, 1999). This is a new surface normal smoothing process which uses the field of principal curvature directions. Rather than directly smoothing the surface normal vectors, we exert control over their directions by smoothing the field of principal curvature vectors. To do this we develop a topography sensitive smoothing process which overcomes the problems of singularities in the field of principal curvature directions at the locations of umbilics and saddles. The method is evaluated on both synthetic and real-world images.

The observation motivating this approach is that although the use of topographic information in the original Worthington and Hancock method (Worthington and Hancock,

1999) improves the pattern of surface normals, it overlooks a considerable body of information concerning the differential structure of the underlying surface residing in the local Darboux frame. At each point on the surface the z axis of the Darboux frame points in the direction of the local surface normal, while the x and y axes are aligned in the directions of the local maximum and minimum surface curvatures.

Our aim in this chapter is to refine the information returned by SFS, not by smoothing the field of surface normal directions, but by smoothing the field of principal curvature directions. There are two novel contributions. First, to overcome the problems encountered at umbilics (the centres of cups and domes) and saddles, we develop a topography sensitive weighting process. Second, we incorporate the smoothing process into the Worthington and Hancock (Worthington and Hancock, 1999) framework which ensures that the recovered surface normals satisfy the data-closeness constraints imposed by the image irradiance equation. By smoothing in this way we avoid using surface normal directions in parabolic surface regions, i.e. at the locations of ridges and ravines. These topographic features are associated with rapid rates of change in surface normal direction, and hence, rapid variations in the local image intensity (where the image gradient is large). Since the change in principal curvature direction is smaller than the change in surface normal direction at such features, Darboux smoothing may moderate the over-smoothing of ridge and ravine topography. Experiments on synthetic and real-world data reveal that the method delivers surface orientation information that is both accurate and stable.

However, there are several problems that limit these curvature sensitive smoothing methods. The main problems are the effects of foreshortening for highly inclined surfaces. Specifically, the Hessian matrix is not invariant under changes in surface inclination. As a result estimates of the Hessian matrix extracted from the Gauss map (the field of surface normals) of a surface will exhibit biases in slant and tilt angle. Hence, in this chapter we also return to the problem of curvature sensitive smoothing using Woodham's viewpoint-invariant estimator of the Hessian (Woodham, 1994). This estimator may be computed directly from the field of surface normals, and includes corrections for the local surface slant and tilt. We use the principal curvatures delivered by the viewpoint-invariant Hessian

to perform smoothing. A comparison with the results obtained using the uncompensated Hessian reveals that the new method delivers better topographic features on the occluding limbs of curved objects.

6.1 Local Darboux Frames

Once estimates of the surface normal directions are to hand, we can represent the local surface structure using a Darboux frame. To compute the Darboux frame, we commence from the Hessian matrix

$$H = \begin{bmatrix} \frac{\partial}{\partial x}(\vec{N}^k)_x & \frac{\partial}{\partial x}(\vec{N}^k)_y \\ \\ \frac{\partial}{\partial y}(\vec{N}^k)_x & \frac{\partial}{\partial y}(\vec{N}^k)_y \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} \\ \\ h_{21} & h_{22} \end{bmatrix} \tag{6.1}$$

where for simplicity we have suppressed the coordinate suffixes on the surface normals. In practice, we approximate the Hessian matrix using directional first derivatives computed from the current estimates of the surface normals. At the pixel with row and column indices (i, j), the estimate of the Hessian is

$$H = \begin{bmatrix} \frac{(\vec{N}^k_{i+1,j})_x - (\vec{N}^k_{i,j})_x}{(x_{i+1,j}) - (x_{i,j})} & \frac{(\vec{N}^k_{i+1,j})_y - (\vec{N}^k_{i,j})_y}{(y_{i+1,j}) - (y_{i,j})} \\ \\ \frac{(\vec{N}^k_{i,j+1})_x - (\vec{N}^k_{i,j})_x}{(x_{i,j+1}) - (x_{i,j})} & \frac{(\vec{N}^k_{i,j+1})_y - (\vec{N}^k_{i,j})_y}{(y_{i,j+1}) - (y_{i,j})} \end{bmatrix} \tag{6.2}$$

where $(x_{i,j}, y_{i,j})$ are the coordinates of the pixel (i, j). The two eigenvalues of the Hessian matrix are the maximum and minimum curvatures:

$$\begin{bmatrix} \lambda^k_M &=& -\frac{1}{2}(h_{11} + h_{22} + S) \\ \lambda^k_m &=& -\frac{1}{2}(h_{11} + h_{22} - S) \end{bmatrix} \tag{6.3}$$

where $S^2 = (h_{11} - h_{22})^2 + 4(h_{12}h_{21})$.

The integrability condition ($h_{12} = h_{21}$) is frequently used as an additional smooth-

ness constraint in SFS to render the recovery of surface height tractable (Frankot and Chellappa, 1988). The estimate of the Hessian given in Equation (6.1) does not in general result in values of the quantities h_{12} and h_{21} which are identical. Hence, we do not enforce integrability at every point on the surface. However, if the quantity S^2 is negative, i.e. S is undefined, then surface integrability is enforced using the equation $h_{12} = h_{21} = (h_{12} + h_{21})/2$ which results in a positive value for $h_{12}h_{21}$ and hence a real value for S. The two curvatures λ_M^k and λ_m^k are hence guaranteed to be real-valued.

On the tangent plane to the surface, the eigenvectors associated with the two curvatures are given by the two-component vectors

$$\vec{M}_2^k = \begin{cases} (\frac{1}{2}(h_{11} - h_{22} + S), h_{21})^T & , \quad h_{11} \geq h_{22} \\ (-h_{12}, \frac{1}{2}(h_{11} - h_{22} - S))^T & , \quad h_{11} < h_{22} \end{cases} \tag{6.4}$$

$$\vec{m}_2^k = \begin{cases} (h_{12}, -\frac{1}{2}(h_{11} - h_{22} + S))^T & , \quad h_{11} \geq h_{22} \\ (\frac{1}{2}(h_{11} - h_{22} - S), h_{21})^T & , \quad h_{11} < h_{22} \end{cases} \tag{6.5}$$

We compute the eigenvector associated with the maximum curvature. This is referred to as the principal curvature direction. From the known direction of the surface normal $\vec{N}^k$ and the principal curvature direction on the tangent plane $\vec{M}_2^k$ we compute the three-component unit-vector $\vec{M}^k$ in the image coordinate system. The direction of the surface normal together with the principal curvature direction define a local frame of reference known as a Darboux frame.

From the two eigenvalues of the Hessian, we can also compute the shape-index of Koenderink and van Doorn (Koenderink and van Doorn, 1992). The quantity is angular in nature and is defined to be

$$\phi^k = \frac{2}{\pi} \arctan\left(\frac{\lambda_m^k + \lambda_M^k}{\lambda_m^k - \lambda_M^k}\right) \quad , \quad \lambda_M^k \geq \lambda_m^k \tag{6.6}$$

The shape-index varies smoothly in the range between -1 and $+1$ as the topography of the surface changes. The shape-index does not depend on the magnitude of the surface

curvature. To capture the degree of bending or magnitude of the curvature, the curvedness quantity $K = \sqrt{(\lambda_m^k)^2 + (\lambda_M^k)^2}$ is used. In our experiments, we use this measure to visualize the whereabouts of curved surface features.

The relationship between the shape-index and the topographic class of the underlying surface is summarized in Table 6.1. The table lists the topographic classes (i.e. dome, ridge, saddle ridge etc.) and the corresponding shape-index intervals. Later on, we visualize surface topography by mapping the shape-index to a grey-scale value. As we move from top to bottom in the table, then the associated grey-scale value decreases.

Class	Region-type	Shape index
Dome	Elliptic	$\left[\frac{5}{8}, 1\right)$
Ridge	Parabolic	$\left[\frac{3}{8}, \frac{5}{8}\right)$
Saddle ridge	Hyperbolic	$\left[\frac{1}{8}, \frac{3}{8}\right)$
Plane	Hyperbolic	Undefined
Saddle-point	Hyperbolic	$\left[-\frac{1}{8}, \frac{1}{8}\right)$
Saddle-rut	Hyperbolic	$\left[-\frac{3}{8}, -\frac{1}{8}\right)$
Rut	Parabolic	$\left[-\frac{5}{8}, -\frac{3}{8}\right)$
Cup	Elliptic	$\left[-\frac{5}{8}, -1\right)$

Table 6.1: Topographic classes and the corresponding shape-index intervals.

6.2 Viewpoint-Invariant Curvatures

The Hessian matrix H can also be represented using the gradient of the surface (p, q), in the following manner.

$$H = \begin{bmatrix} \frac{\partial p}{\partial x} & \frac{\partial p}{\partial y} \\ \frac{\partial q}{\partial x} & \frac{\partial q}{\partial y} \end{bmatrix} \tag{6.7}$$

For integrable surfaces H is symmetric, i.e. $\frac{\partial p}{\partial y} = \frac{\partial q}{\partial x}$. The Hessian H is a viewer-centered representation of surface curvature and hence depends on the orientation of the surface. From the Hessian H and the gradient (p, q) we can compute the viewpoint-invariant representation of surface curvature developed by Woodham (Woodham, 1994)

and given by

$$C = \frac{1}{\sqrt{(1+p^2+q^2)^3}} \begin{bmatrix} q^2+1 & -pq \\ -pq & p^2+1 \end{bmatrix} \begin{bmatrix} h_{11} & h_{12} \\ h_{21} & h_{22} \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{bmatrix} \tag{6.8}$$

The two eigenvalues of this viewpoint-compensated Hessian matrix C are the maximum and minimum curvatures. These curvature values are viewpoint-invariant surface properties because they do not depend on the viewer-centered coordinate system. The terms in Equation (6.8) involving the gradient can be interpreted as the corrections necessary to account for the geometric foreshortening associated with viewing a surface element obliquely. The maximum and minimum curvature directions obtained using the matrix C are viewpoint dependent, i.e. although they are orthogonal in the object-centered coordinate system defined by the local surface normal and tangent plane, they are not, in general, orthogonal when projected onto the image plane (Woodham, 1994).

6.3 Smoothing Principal Curvature Directions

We would like to smooth the field of surface normal directions using the field of principal curvature directions. However, this is not a straightforward task. The reason for this is that certain classes of local surface topography are characterized by rapid changes in principal curvature direction. There are two important cases. Firstly, for elliptic structures, i.e. domes and cups, there is a singularity since the principal curvature is not defined at the associated umbilic or spherical point, and the surrounding vector field is radial in structure. Secondly, for hyperbolic structures, i.e. saddles, there is a discontinuity in the principal curvature direction. This is due to the fact that the directions of maximum and minimum curvature flip as the saddle is traversed.

To overcome these problems we weight against domes and saddles in the smoothing process. To do this we note that the shape-index associated with parabolic structures is

± 0.5. We therefore use the topographic weight

$$w_{i,j}^{(k)} = \exp\{-\mu(|\phi_{i,j}^k| - 0.5)^2\} \tag{6.9}$$

where μ is a constant. The weight is used to compute the smoothed principal curvature direction using the formula

$$\vec{M}_S^k(i,j) = \sum_{(l,m)\in R_{i,j}} w_{l,m}^{(k)} \vec{M}^k(l,m) \tag{6.10}$$

over the neighborhood $R_{i,j}$ of the pixel indexed (i,j). Our aim is to use the smoothed principal curvature direction to compute an updated Darboux frame.

6.4 Surface Normal Adjustment

Once the smoothing of the field of principal curvature directions is complete, then we can use their direction vectors to compute revised estimates of the surface normal directions ($\vec{N}_D^k$). This is a two-step process. First, we rotate the unsmoothed surface normals around the irradiance cone in a manner which is consistent with the rotation of the principal curvature directions which results from the Darboux smoothing process. To do this, we project the smoothed and unsmoothed principal curvature directions onto the plane which is perpendicular to the axis of the irradiance cone (i.e. the light-source direction). Once these projections are to hand, then we compute their angle-difference on the plane. We rotate the on-cone surface normal around the cone by an amount equal to this angle. The second step is to restore the orthogonality of the Darboux frame by rotating the surface normal away from the irradiance cone so that it is perpendicular to the smoothed principal curvature direction. The procedure adopted is illustrated in Figures 6.1 and 6.2.

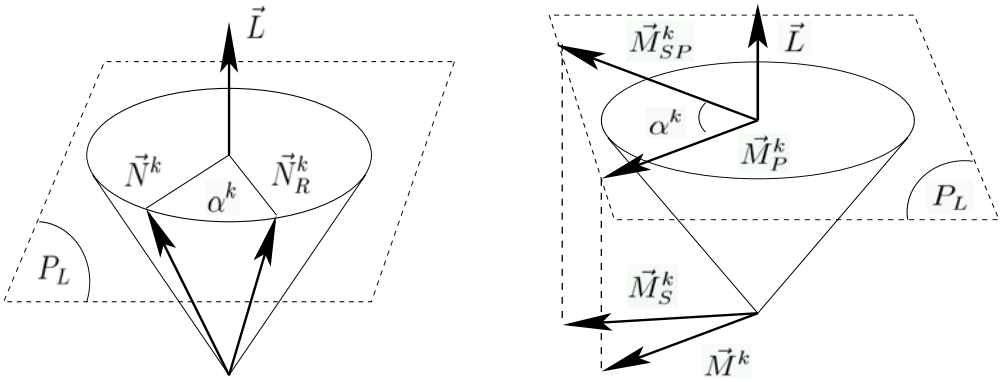

Figure 6.1: Smoothing the local Darboux frame: (a) rotating the current surface normal around the irradiance cone by the angle α^k; (b) computing the rotation angle α^k using the unsmoothed and smoothed principal curvature directions, and their projections.

6.4.1 Rotation About the Irradiance Cone

To be more formal, the projection of the unsmoothed principal curvature direction onto the plane P_L perpendicular to the axis of the Lambertian irradiance cone is

$$\vec{M}_P^k = \vec{L} \times \vec{M}^k \times \vec{L} \tag{6.11}$$

while the projection of the smoothed principal curvature direction onto the plane P_L is

$$\vec{M}_{SP}^k = \vec{L} \times \vec{M}_S^k \times \vec{L} \tag{6.12}$$

On the plane P_L, the unsigned angle between the unit vectors of the projected principal curvature directions, i.e. $\vec{M}_P^k$ and $\vec{M}_{SP}^k$, is

$$|\alpha^k| = \cos^{-1}(\vec{M}_P^k \cdot \vec{M}_{SP}^k) \tag{6.13}$$

To perform the rotation of the on-cone surface normal in a consistent manner, we require the sign or sense of this angle. Because the vectors $\vec{M}_P^k$ and $\vec{M}_{SP}^k$ are both perpendicular to the light-source direction $\vec{L}$, the cross product $\vec{M}_P^k \times \vec{M}_{SP}^k$ is either in the direction $\vec{L}$ or in the opposite direction. So, we can compute the signed rotation angle α^k

using the formula

$$\alpha^k = \overbrace{(\vec{L} \cdot (\vec{M}_P^k \times \vec{M}_{SP}^k))}^{sign} |\alpha^k| \tag{6.14}$$

We rotate the current surface normal around the irradiance cone, defined by the light-source direction $\vec{L}$ and the apex angle $\cos^{-1}(E)$, by the angle α^k. The rotation is illustrated in Figure 6.1. Hence, the updated surface normal direction is given by the formula

$$\vec{N}_R^k = \vec{N}^k \cos(\alpha^k) + (\vec{N}^k \cdot \vec{L})\vec{L}(1 - \cos(\alpha^k)) + (\vec{L} \times \vec{N}^k) \sin(\alpha^k) \tag{6.15}$$

6.4.2 Computing the Darboux Surface Normal Direction

The surface normals computed using Equation (6-15) are no longer perpendicular to the two principal curvature directions. We must therefore provide a means of restoring the geometry of the Darboux frame.

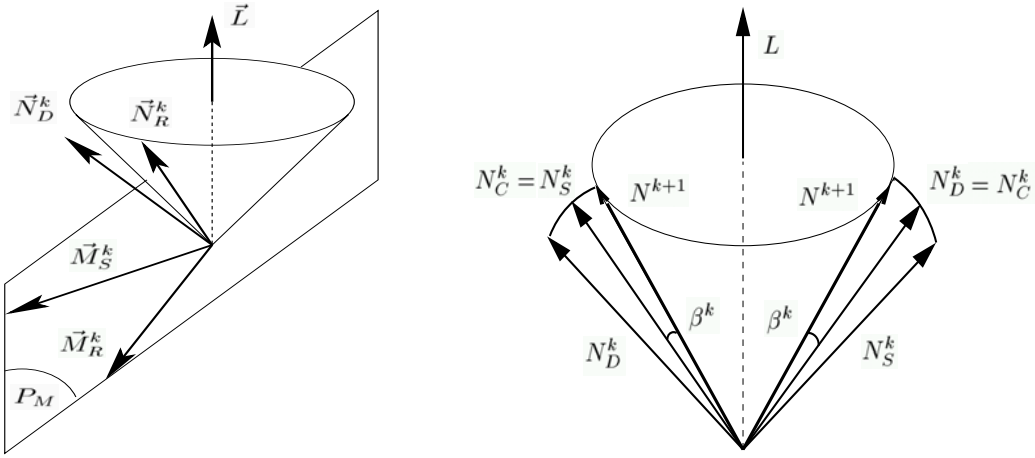

Figure 6.2: Computing the updated surface normal direction: (a) computing the direction of local surface normal to the smoothed Darboux frame; (b) rotating the closest surface normal onto the irradiance cone to impose the irradiance constraint.

To do this, we use the field of rotated surface normals $\vec{N}_R^k$ to compute an estimate of the principal curvature direction, which we denote by $\vec{M}_R^k$ (Figure 6.2.a). The vector that is perpendicular to $\vec{M}_R^k$ and the smoothed principal curvature direction $\vec{M}_S^k$ provides an average surface normal. However, we must test the sign of the z-component of the

resulting vector to ensure that the handedness of the coordinate system does not flip. The surface normal used is

$$\vec{N}_D^k = \begin{cases} \vec{M}_S^k \times \vec{M}_R^k & , \quad (\vec{M}_S^k \times \vec{M}_R^k)_z \geq 0 \\ \vec{M}_R^k \times \vec{M}_S^k & , \quad (\vec{M}_S^k \times \vec{M}_R^k)_z < 0 \end{cases} \tag{6.16}$$

6.4.3 Imposing the Irradiance Constraint

To restore the irradiance constraint we can clearly apply the procedure of Worthington and Hancock and project the normals onto the nearest position on the irradiance cone. However, there are situations where the principal curvature directions are unavailable. For instance, the principal curvature direction is not defined at the location of umbilics. Also, it can not be computed when the Hessian matrix is singular.

There is a further reason to use smoothed surface normals $\vec{N}_S^k$ in conjunction with Darboux surface normals $\vec{N}_D^k$. The reason for this is that there are certain surface regions for which the Darboux surface normals $\vec{N}_D^k$ may be more accurate and may be closer to the irradiance cone than the smoothed surface normals $\vec{N}_S^k$. The cases in question are a) locations where image intensity gradients are large; b) at the locations of ridges and ruts. On the other hand, the smoothed surface normals $\vec{N}_S^k$ may be more accurate a) where the intensity gradients are small and b) at the locations of domes.

To overcome these problems, we have used a switching process. For each location in the image we have computed two surface normal directions. The first of these is the surface normal $\vec{N}_D^k$ computed from the smoothed principal curvature direction. The second is the surface normal $\vec{N}_S^k$ found by applying the robust smoothing method of Worthington and Hancock to the field of current surface normal estimates $\vec{N}^k$. From these two alternatives, we select the surface normal which is closest to the irradiance cone. In other words, the surface normal used for projection onto the irradiance cone is

$$\vec{N}_C^k = \begin{cases} \vec{N}_D^k & , \quad |\cos^{-1}(\vec{N}_D^k \cdot \vec{L}) - \cos^{-1}(E)| \leq |\cos^{-1}(\vec{N}_S^k \cdot \vec{L}) - \cos^{-1}(E)| \\ \vec{N}_S^k & , \quad |\cos^{-1}(\vec{N}_D^k \cdot \vec{L}) - \cos^{-1}(E)| > |\cos^{-1}(\vec{N}_S^k \cdot \vec{L}) - \cos^{-1}(E)| \end{cases} \tag{6.17}$$

With the closest surface normal to hand, we can effect projection onto the irradiance cone. The rotation angle required for the projection is

$$\beta^k = \cos^{-1}(\vec{N}_C^k \cdot \vec{L}) - \cos^{-1}(E) \tag{6.18}$$

The geometry of this procedure is described in Figure 6.2.b. Hence, after projection the updated surface normal is

$$\vec{N}^{k+1} = \vec{N}_C^k \cos(\beta^k) + \left(\left(\frac{\vec{N}_C^k \times \vec{L}}{\|\vec{N}_C^k \times \vec{L}\|} \right) \times \vec{N}_C^k \right) \sin(\beta^k) \tag{6.19}$$

6.5 Experiments

We have experimented with the Darboux smoothing method on synthetic and real-world imagery. The experimentation with synthetic data is designed to measure the effectiveness of the method on surfaces which contain a variety of topographic structure. Since our method relies on well defined principal curvature directions, we may anticipate that the main advantages of our method are most noticeable at ridge or ravine structures. We have therefore conducted some experiments on real and synthetic data aimed at evaluating our Darboux smoothing (D-S) process at such structures and comparing its performance with the Worthington and Hancock (W-H) method (Worthington and Hancock, 1999).

To visualize the shape-index ϕ (computed using Equation 6.6) for the field of surface normal directions obtained using SFS, the quantity needs to be rescaled from the interval $[-1, +1]$ to the grey-scale interval $[0, 255]$. Hence, we display the quantity $\frac{255(1+\phi)}{2}$ at each image location. To aid visualization, we also scale the curvedness values (which are always positive) to fall in the grey-scale interval $[0, 255]$.

We have compared our D-S method with the W-H method (Worthington and Hancock, 1999). Their method does not make use of the principal curvature directions. However, the principal curvature directions corresponding to the surface normal directions obtained using their method can easily be computed using the method outlined in Section 6.1.

In our experiments, we provide comparison by computing the difference in surface normal directions. These are displayed as fields of differences in unit surface normal directions. The vectors in the resulting vector fields show the direction and relative magnitude of the vector difference in the surface normals. We also visualize the differences in surface normal directions by displaying the dot-product between the unit surface normals (this quantity takes on it maximum value of unity when the directions of the surface normal are identical). The dot-products are mapped onto the grey-scale interval, and hence the brighter the pixel, the smaller the directional error.

6.5.1 Synthetic Data

Figures 6.3-6.7 show the results obtained with synthetic surfaces. In Figures 6.3-6.5, in the top row, we show the elevation data for the surface (left panel) and a rendering of the local surface orientation information using Lambert's law (right panel). The second, third and fourth rows respectively show ground-truth (G-T), the results obtained with the W-H method, and, those obtained with our D-S method. In these three rows, the first (left-hand) column shows the surface normal data, the second column shows the field of minimum curvature directions, the third column is the total surface shape-index.

The surface studied in Figure 6.3 is obtained by placing four elliptical domes on a parabolic ridge. In Figure 6.4 the surface consists of three narrow paraboloid ridges placed on top of perpendicular wide ridge. The surface in Figure 6.5 is similar to that in Figure 6.4, except that the four surfaces have their maxima at the same height.

For the surface shown in Figure 6.3, the minimum curvature vectors radiate from the four domes, and are parallel to the crest line of the parabolic ridge. There should be a sharp discontinuity in direction where the domes intersect the ridge. Both smoothing methods blur this transition region. However, in the case of our D-S method the blurring is less pronounced. This feature emerges more clearly in the shape-index images, where the blurring associated with the W-H method is very strong. By contrast the D-S method produces results which are in good agreement with ground-truth.

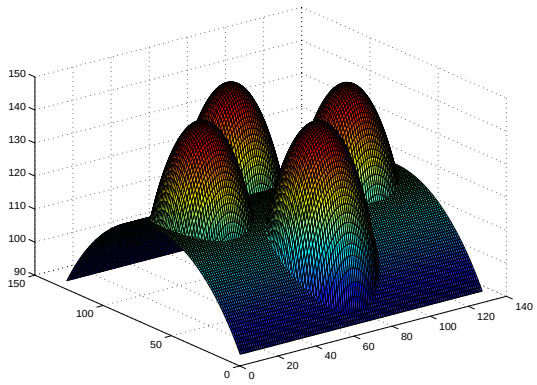

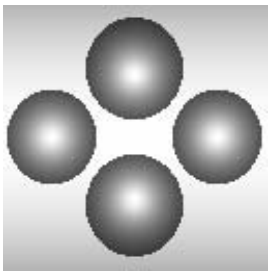

brightness data

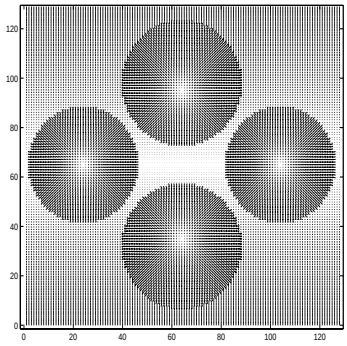

surface normals (G-T)

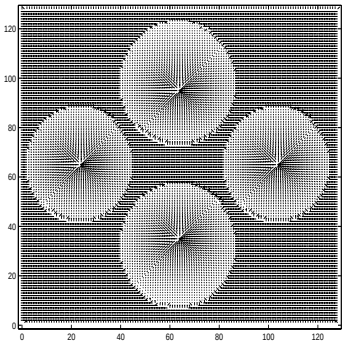

minimum curvature dir. (G-T)

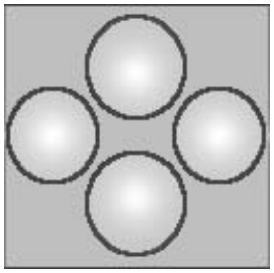

shape-index (G-T)

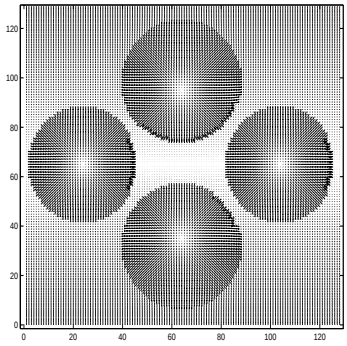

surface normals (W-H)

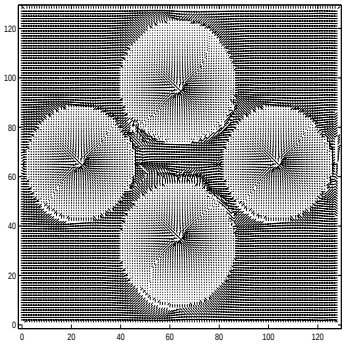

minimum curvature dir. (W-H)

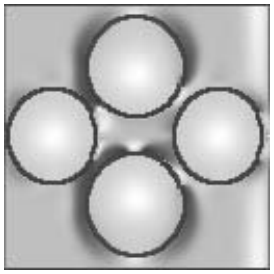

shape-index (W-H)

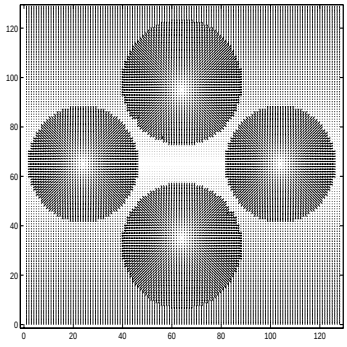

surface normals (D-S)

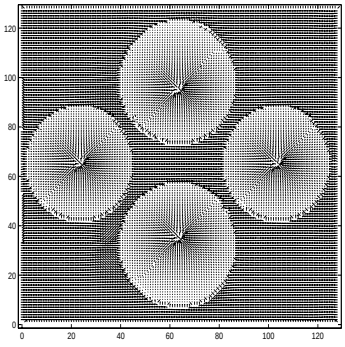

minimum curvature dir. (D-S)

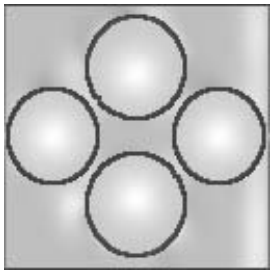

shape-index (D-S)

Figure 6.3: The paraboloid-ellipsoids: ground-truth data (rows 1-2); Worthington-Hancock method (row 3); Darboux smoothing method (row 4).

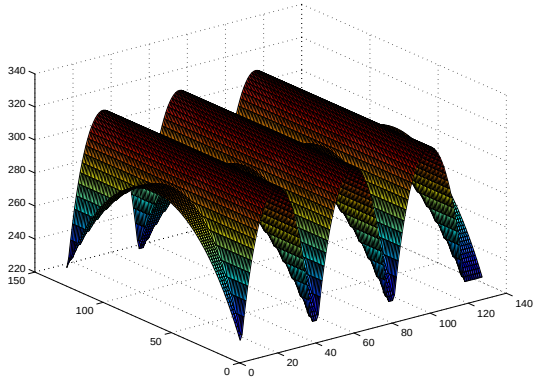

brightness data

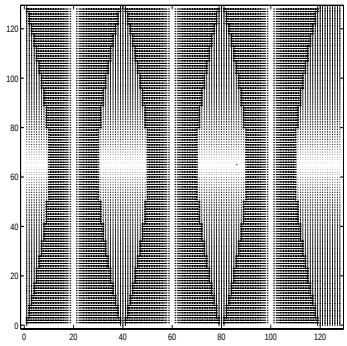

surface normals (G-T)

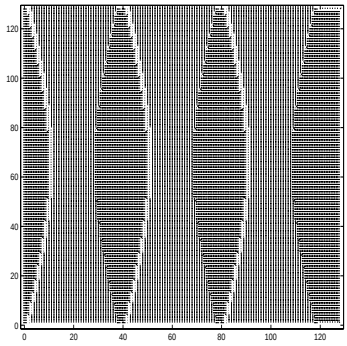

minimum curvature dir. (G-T)

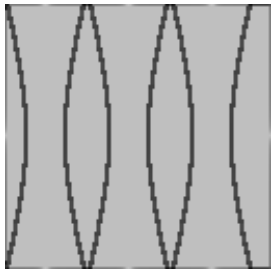

shape-index (G-T)

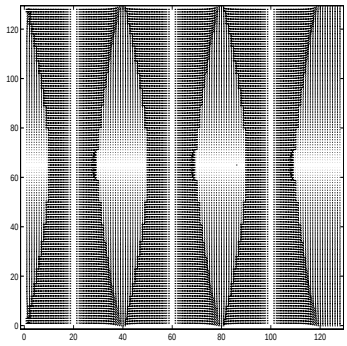

surface normals (W-H)

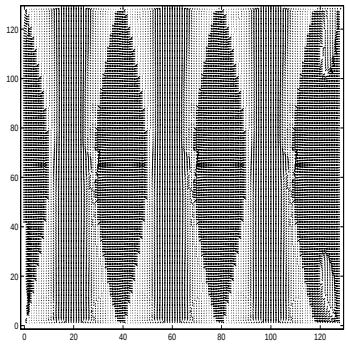

minimum curvature dir. (W-H)

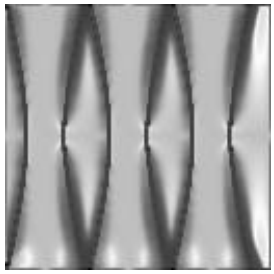

shape-index (W-H)

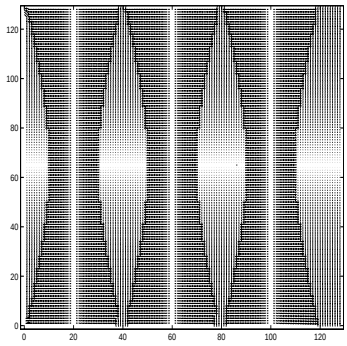

surface normals (D-S)

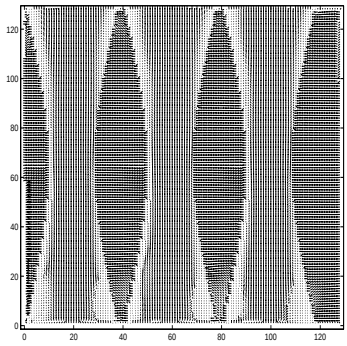

minimum curvature dir. (D-S)

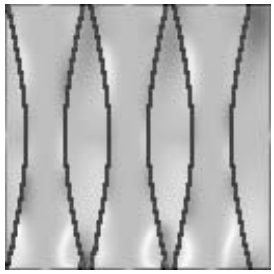

shape-index (D-S)

Figure 6.4: The crossed paraboloids with two levels of maxima: ground-truth data (rows 1-2); Worthington-Hancock method (row 3); Darboux smoothing method (row 4).

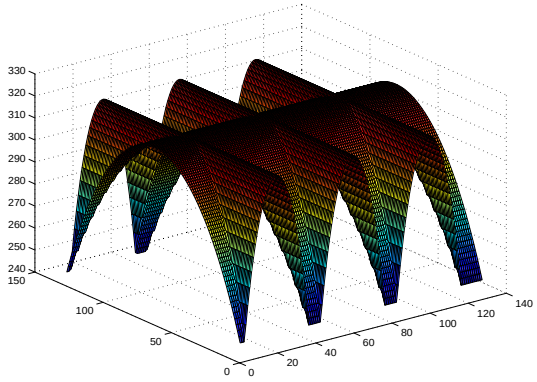

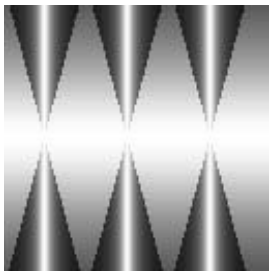

brightness data

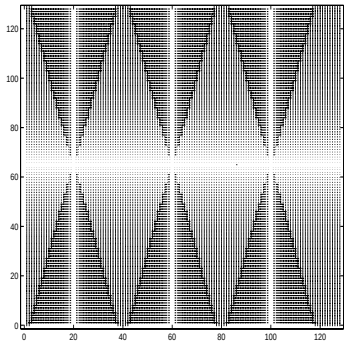

surface normals (G-T)

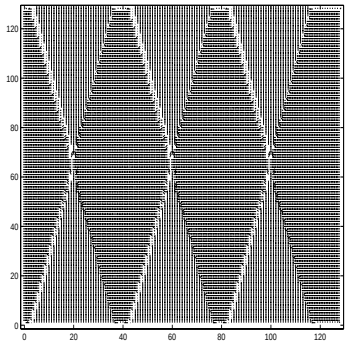

minimum curvature dir. (G-T)

shape-index (G-T)

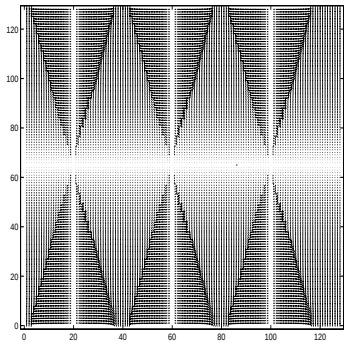

surface normals (W-H)

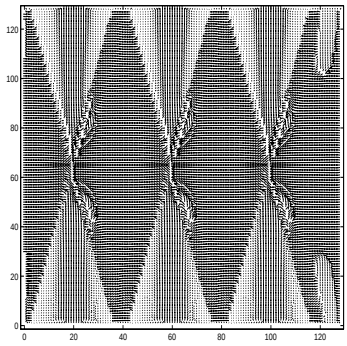

minimum curvature dir. (W-H)

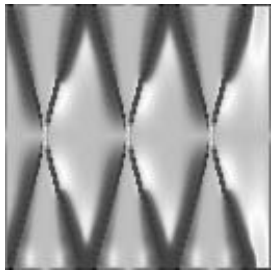

shape-index (W-H)

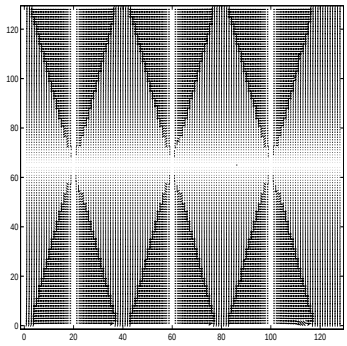

surface normals (D-S)

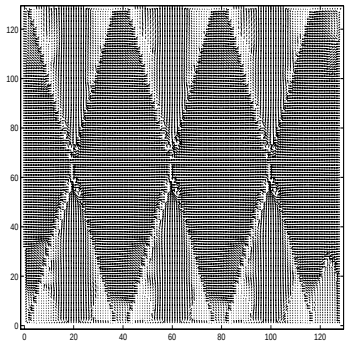

minimum curvature dir. (D-S)

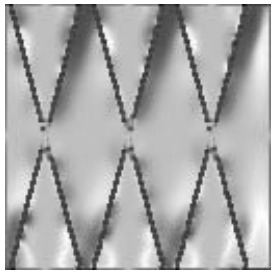

shape-index (D-S)

Figure 6.5: The crossed paraboloids with identical levels of maxima: ground-truth data (rows 1-2); Worthington-Hancock method (row 3); Darboux smoothing method (row 4).

In Figures 6.4 and 6.5, the pattern of minimum curvature directions is more straightforward. In each of the ridge surfaces the minimum curvature direction is parallel to the crest line. Hence, for the narrow ridges, the minimum curvature direction points in the y (vertical) direction, while for the wide ridge, it points in the x (horizontal) direction. Again, the D-S method gives better agreement with ground-truth and results in less blurring at the intersection of the different surfaces.

In Figure 6.6, the columns from left to right are for the three synthetic surfaces used in Figures 6.3, 6.4 and 6.5, respectively. The first and the second rows show the field of difference vectors between the ground-truth and the surface normals obtained using the W-H method (first row) and our D-S method (second row). Furthermore, the third and the fourth rows show the images of dot-products between the estimated and the ground-truth surface normal directions. For the first column, from the difference fields and dot-product panels it is clear that the main departures between the two smoothing methods occur in the dome regions of the surface.

The plots in Figure 6.7 present some data to illustrate the iterative qualities of the two algorithms. The curves show the total-mean-square (TMS) error between the ground-truth and the estimated surface normals as a function of iteration number (as defined in Section 4.3.1). Here the solid curve shows the error obtained with D-S, while the dashed curve is the result obtained with the W-H method. The D-S method converges monotonically to a lower value of the TMS error in fewer iterations. In the case of the W-H method, the method initially converges, but then the error rapidly increases with iteration number. This convergence problem with the W-H method is overcome by terminating the method at the optimal number of iterations. When our method is run with this optimal number of iterations, then the error is comparable, and in some cases better than, the W-H method.

Finally, we discuss the additional computational overheads introduced by our D-S method when compared with that of W-H. At each iteration, our method involves the additional computation of the eigenvalues and eigenvectors of the Hessian matrix. This is straightforward since the Hessian matrix is 2x2 and, as shown in Section 6.1, there are simple closed-form expressions for the quantities. In addition, there is the computation of

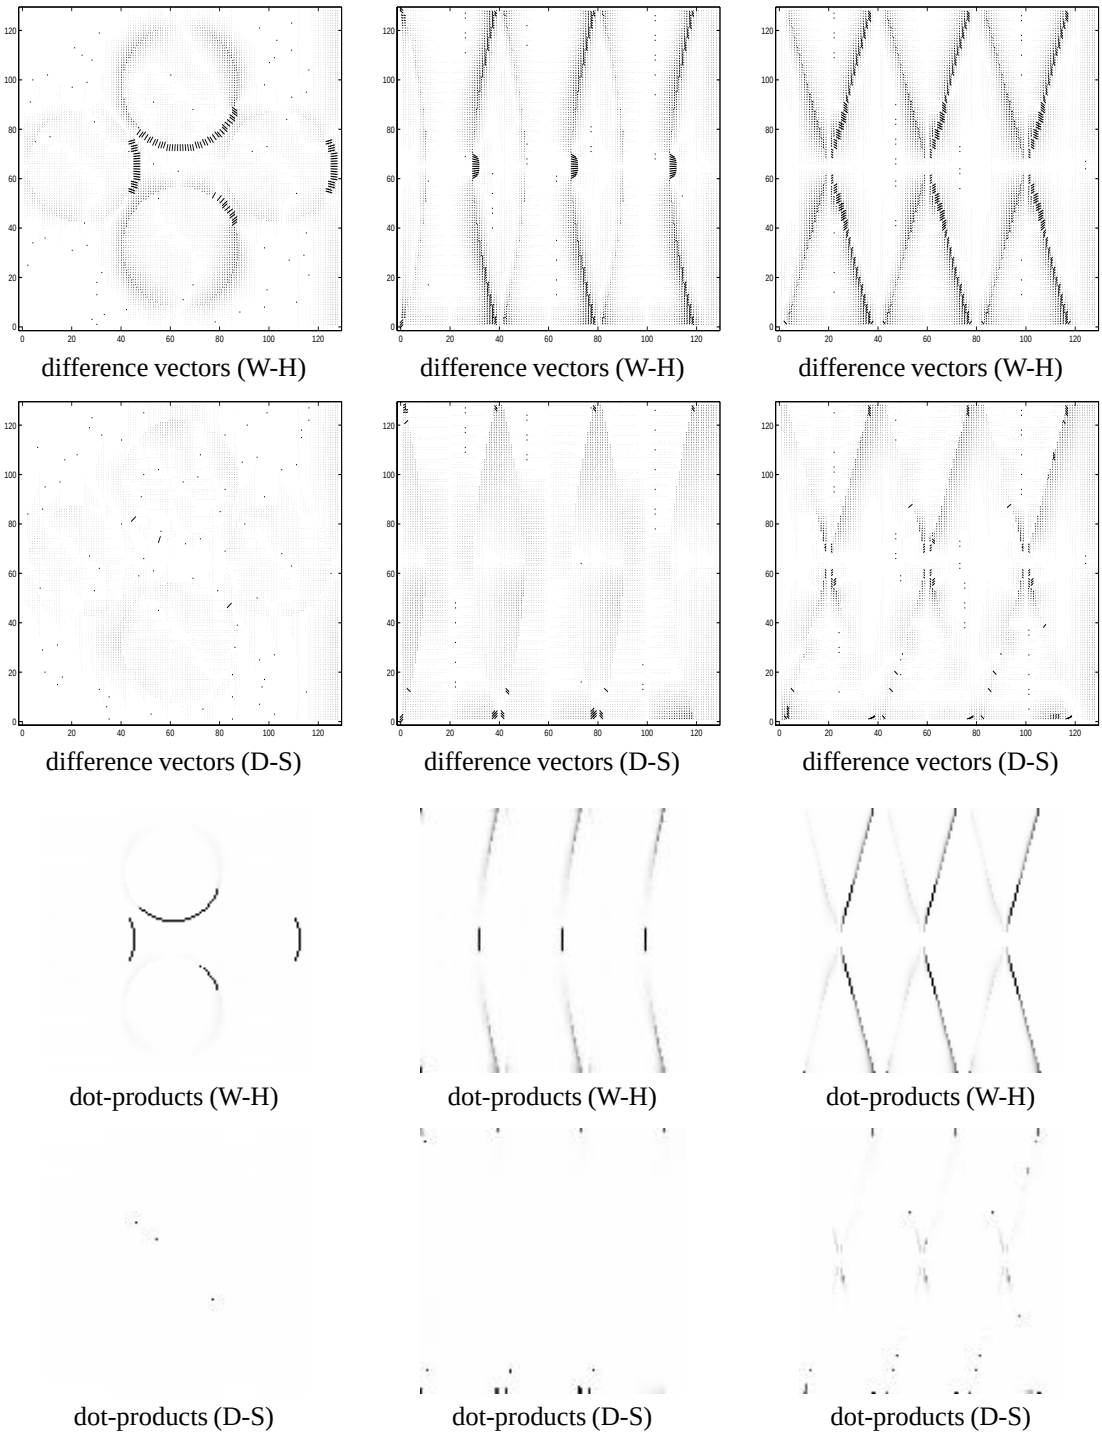

Figure 6.6: Results for the images of Figures 6.3-6.5: difference vectors between the ground-truth surface normals and those obtained using the Worthington-Hancock method and the Darboux smoothing method (rows 1-2); dot-product images (rows 3-4).

the weighted mean principal curvature direction and the rotation of the Darboux frame. Since the amount of D-S depends on the topographic complexity of the surfaces under

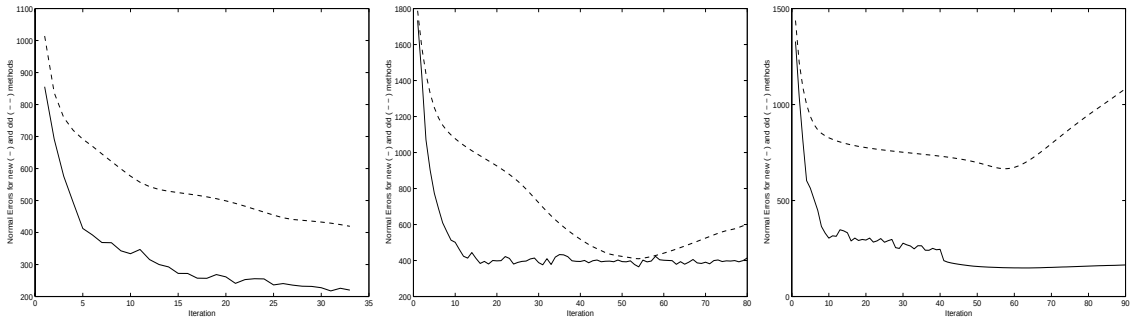

Figure 6.7: TMS surface normal error versus iteration number obtained using the methods of Worthington-Hancock (dashed curves) and Darboux smoothing (solid curves); panels from left to right correspond to the surfaces shown in Figures 6.3, 6.4 and 6.5 respectively.

study, the additional computation required is image dependant. The main effect here is that the number of iterations required increases with the complexity of the topographic structure present. However, as an example for the image shown in Figure 6.3 which is of size 128x128 pixels, the W-H method took 7 secs, while our method took 11 secs (both in 33 iterations) on a Pentium III (700 MHz) machine.

6.5.2 Real-World Data

We have studied a variety of real-world images. The first of these is the toy duck (Figure 6.8) which comes from the Columbia University Object Library (COIL-20) database (Nene et al., 1996). This database contains a number of images of relatively complex shapes viewed from different directions with known light-source direction. We have also captured images of terra-cotta objects (Figures 6.9-6.10) using an Olympus 10E camera. These objects are made of clay and although their surfaces are moderately rough (non-Lambertian), no filtering has been applied to them. Each object has been imaged under controlled lighting conditions in a darkroom. The objects have been illuminated using a single collimated tungsten light-source. The light-source direction is recorded at the time the images are captured.

In Figures 6.8-6.11 we study images of a toy duck, a terra-cotta watering-can and a terra-cotta teapot. In Figures 6.8-6.10, the left panel in the top row is the original image.

The results shown in the first row are the estimated field of surface normals obtained using the W-H method and the D-S method. In the second row, from left to right the panels show the following results. In the first panel we show the field of difference vectors between the two sets of surface normals. The second and third panels show the fields of minimum curvature directions corresponding to the fields of surface normals in the top row. Finally, in Figure 6.11, we use the surface normals obtained using the two methods to compute the shape-index images (the first and the second columns), the curvedness images (the third and the fourth columns) and the dot-product images (the fifth column) for the three objects shown in Figures 6.8-6.10.

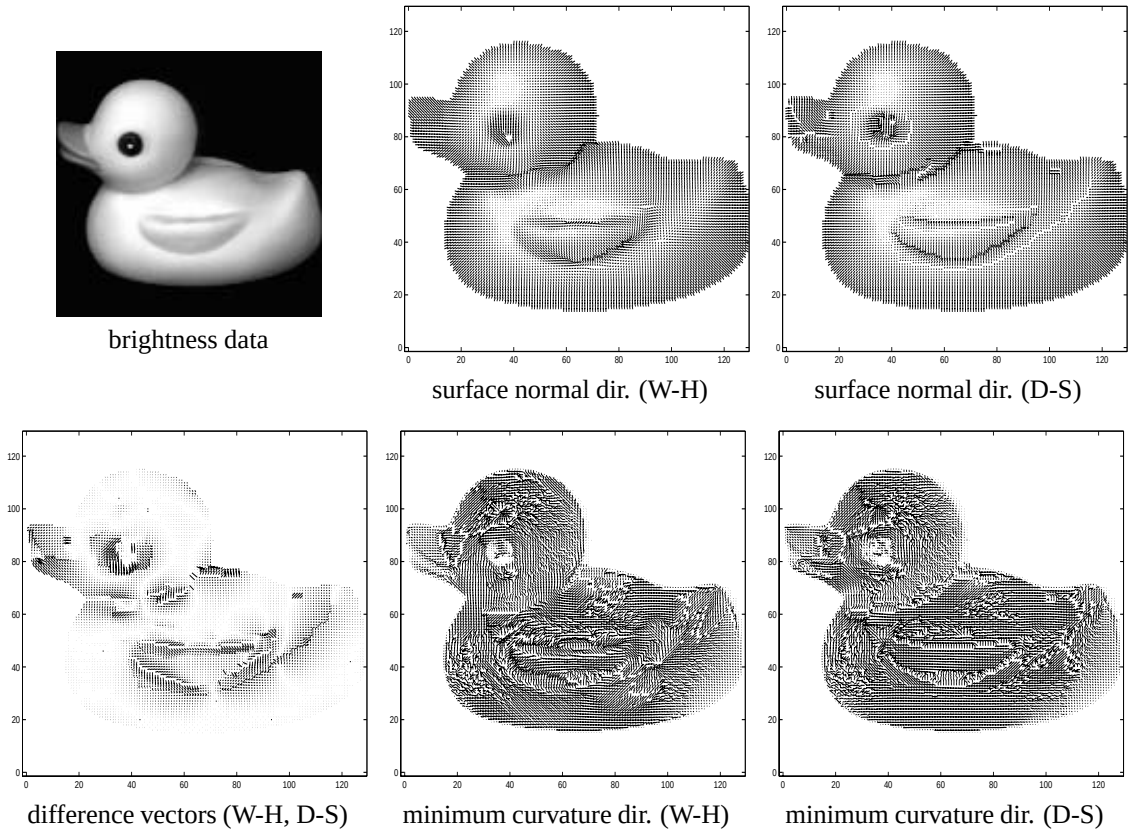

Figure 6.8: A toy duck: surface normals obtained using the two methods (row 1); difference between the two sets of surface normals; minimum curvature directions (row 2).

In the case of the toy duck image (Figure 6.8), there are a number of features to note. First, from the curvedness data, the D-S method locates the ridge line which extends from the tail of the duck to its chest. This structure is not detected by the W-H method. Second,

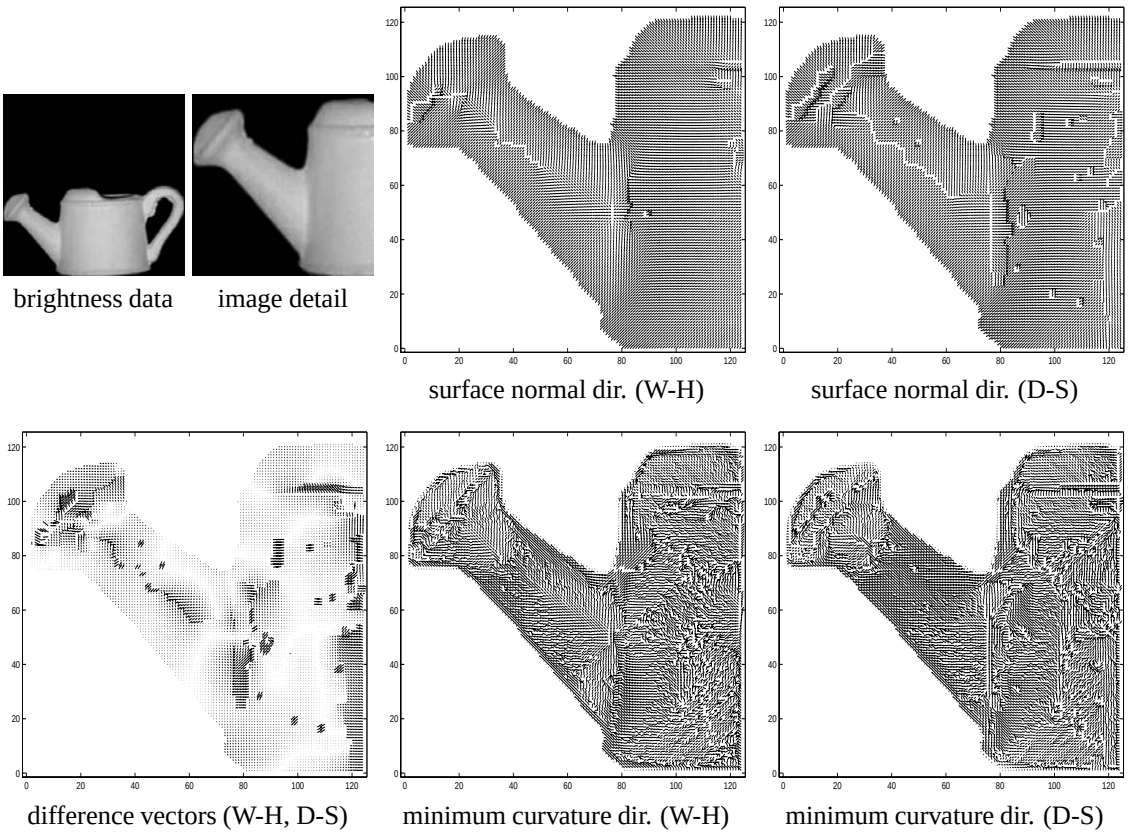

Figure 6.9: A terra-cotta watering-can: surface normals obtained using the two methods (row 1); difference between surface normals; minimum curvature directions (row 2).

from the shape-index data, the D-S method provides cleaner detail around, the neck, wing and beak. These features are also evident in the surface normal difference field.

In the case of the watering-can image (Figure 6.9), there are similar differences. These are mainly associated with the intersection of the spout with the body, and the detail of the spout. From the curvedness data, the D-S method locates the main crest line of the spout, which is a weak feature in the W-H method. In the teapot image (Figure 6.10), we have a similar surface structure around the spout. The D-S method again gives better location of the surface intersection and the structure of the ribbed spout.

6.5.3 Hessian Compensation

In this section we compare the uncompensated and compensated Hessian matrices in the D-S method using a synthetic image and a real-world image. The synthetic data with

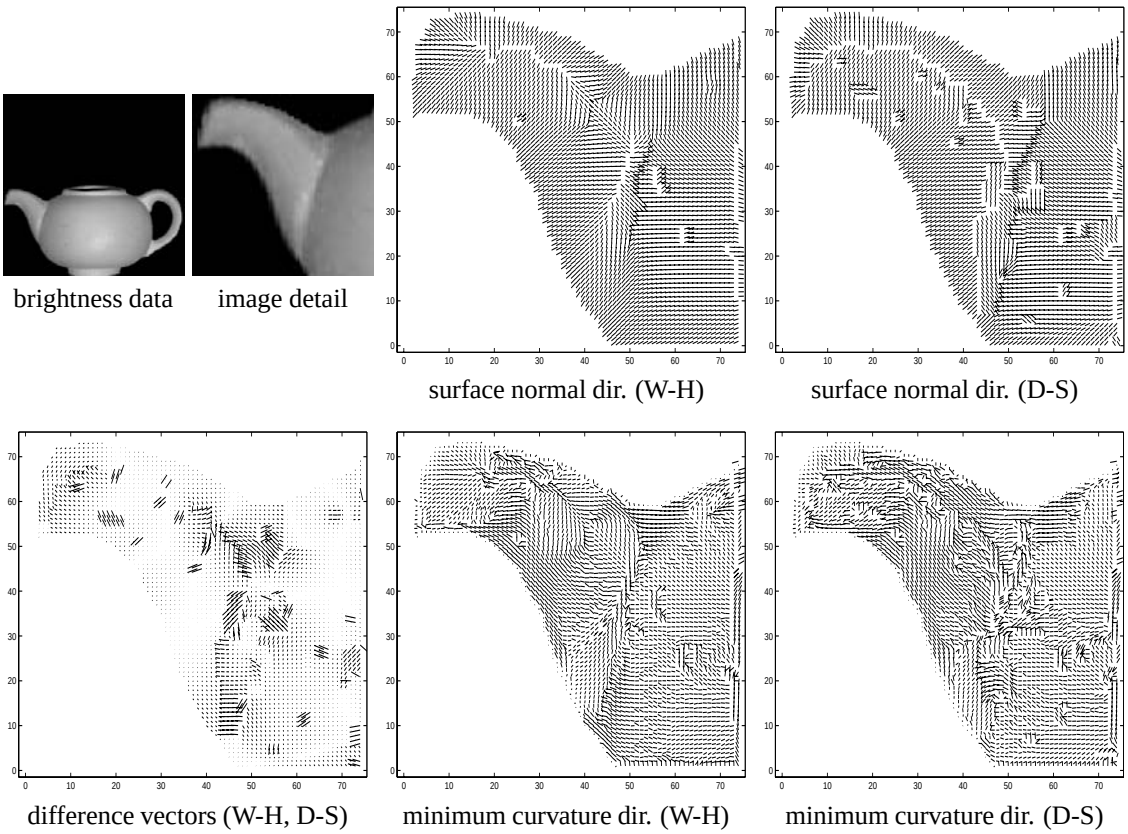

Figure 6.10: A terra-cotta teapot: surface normals obtained using the two methods (row 1); difference between surface normals; minimum curvature directions (row 2).

known ground-truth allows us to explore in more detail the effects of the viewpoint-invariant curvatures outlined in Section 6.2 on the recovered surface normals. The synthetic shape used here consists of a number of small ellipsoids located on top of a large paraboloid. In Figure 6.12, we show the elevation plot of the surface studied together with its synthetic Lambertian image. Figures 6.13-6.16 show the results obtained by applying SFS method using the D-S process to this Lambertian image.

In the top row of Figure 6.13, the left-hand and the right-hand panels show the surface normal directions obtained with the uncompensated Hessian and those obtained with the viewpoint-compensated Hessian, respectively. In the bottom row of the figure, we show the fields of difference between the recovered and the ground-truth surface normals. The main differences for both methods are at the intersections of the outer ellipsoids and the paraboloid. To study the differences more carefully, we show some scatter plots in

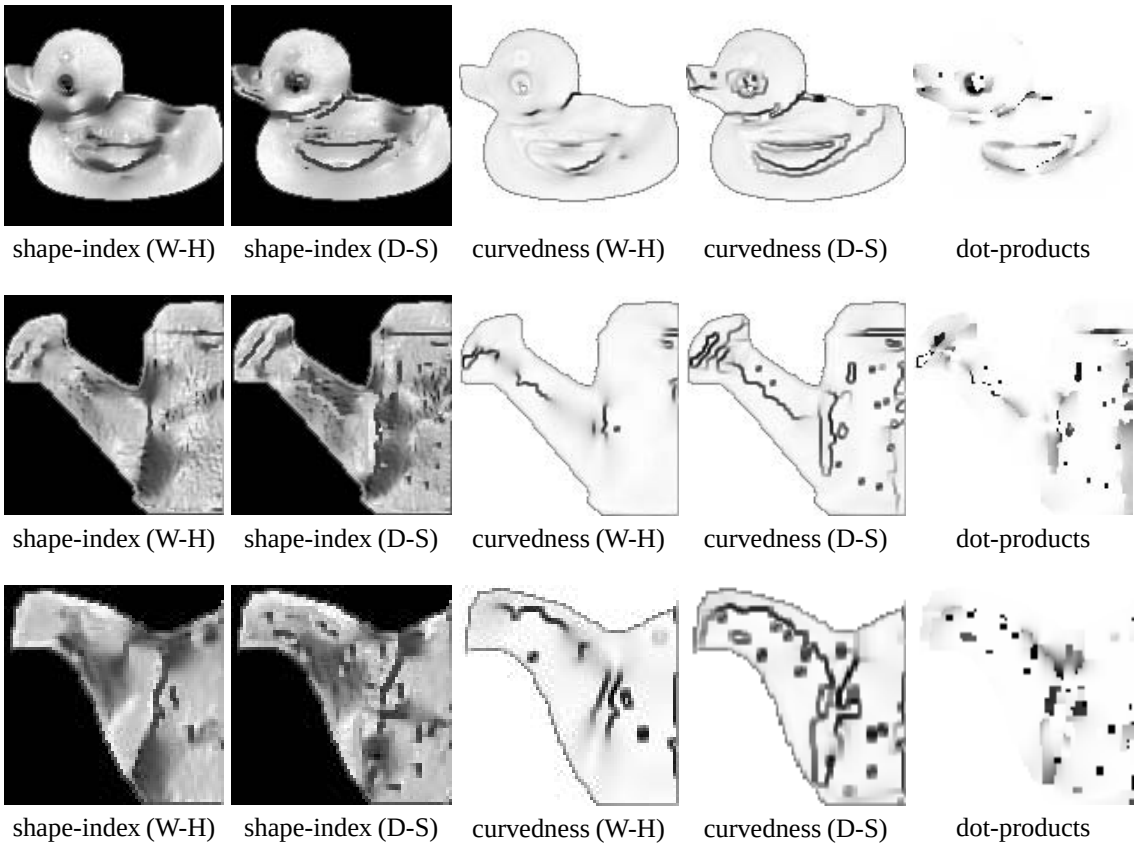

Figure 6.11: Results for the images of Figures 6.8-6.10: (a-b) the shape-index, (c-d) the curvedness and (e) dot-product images for the surface normals obtained using the Worthington-Hancock method and the Darboux smoothing method.

Figure 6.14 of the zenith and azimuth angles for these surface normals. Obviously, the zenith angles, which are the arc cosine of the Lambertian intensity, are constant for all cases, whereas the azimuth angles are variable. In Figure 6.14, the left-hand column is for the uncompensated Hessian whereas the right-hand column is for the viewpoint-compensated Hessian. The plots in the first row show the azimuth angles of the recovered surface normals versus those of the ground-truth data. The plots in the second row show the difference between the recovered and ground-truth azimuth angles versus the zenith angles. It is clear from these scatter plots that the largest differences occur at large zenith angles. Also, the azimuth angles obtained using the compensated Hessian are in general in better agreement with the ground-truth data. This is attributable to the fact that the uncompensated Hessian tends to overestimate the surface curvature at highly inclined locations on the surface due to foreshortening effects. However, there are some locations

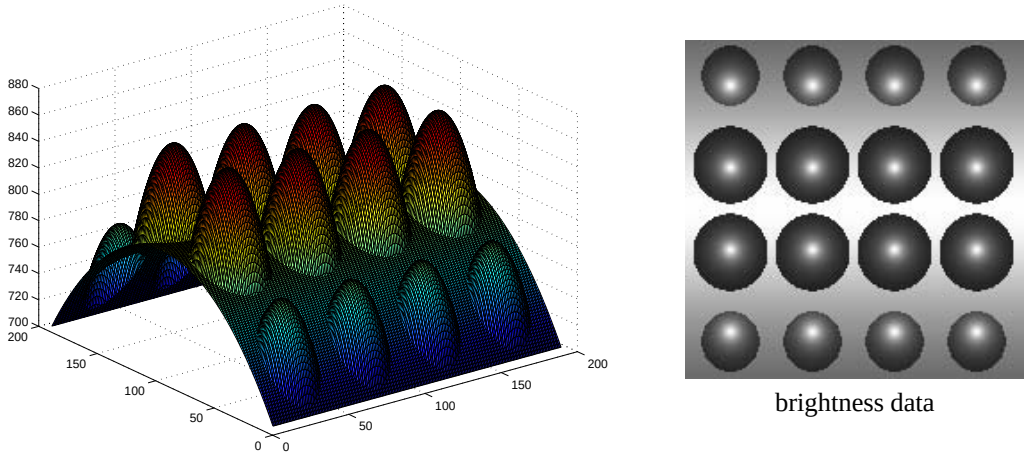

Figure 6.12: A synthetic surface containing 16 ellipsoids on top of a large paraboloid.

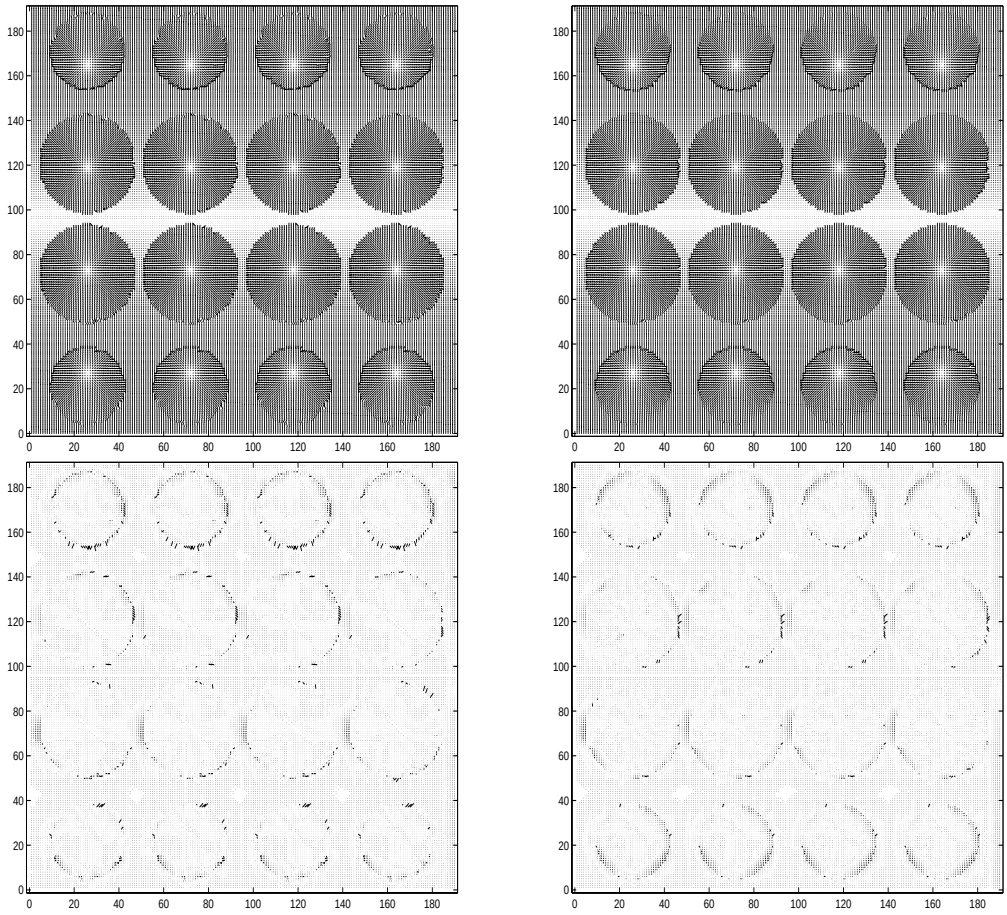

Figure 6.13: Row 1: fields of surface normals using the uncompensated (left) and compensated (right) Hessian in the D-S method; row 2: differences with ground-truth.

at which the uncompensated Hessian results in azimuth angles which are closer to the ground-truth. This is why the viewpoint-invariant Hessian is best used when the highly

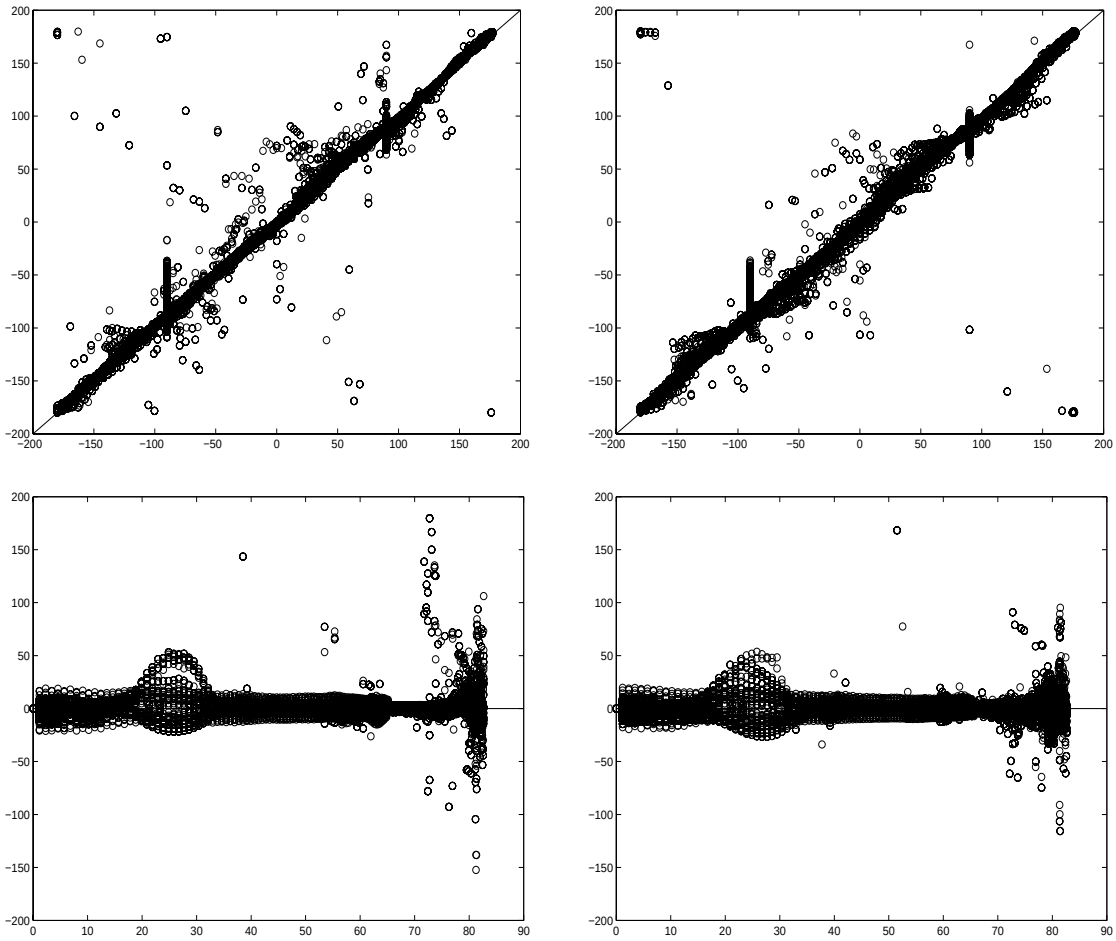

Figure 6.14: Row 1: recovered azimuth angles using the D-S method with the uncompensated (left) and compensated (right) Hessian versus ground-truth azimuth angles; row 2: difference between the recovered and ground-truth azimuth angles versus zenith angles.

inclined locations play an important role.

In Figure 6.15, we compare the shape-index and the curvedness images. The first row is for the shape-index images while the second row is for the curvedness images. From the left to the right, the panels are for the ground-truth, the uncompensated Hessian and the compensated Hessian. Again, it is clear from the shape-index images that the largest differences occur at the intersections of the outer ellipsoids and the paraboloid. Also, the compensated Hessian at these locations results in more consistent shape-index values for similar shapes when compared to the uncompensated Hessian.

In Figure 6.16, we compare the iterative qualities of the two methods. Here, we plot the TMS error at each iteration. Note that we only show the final surface normals (at

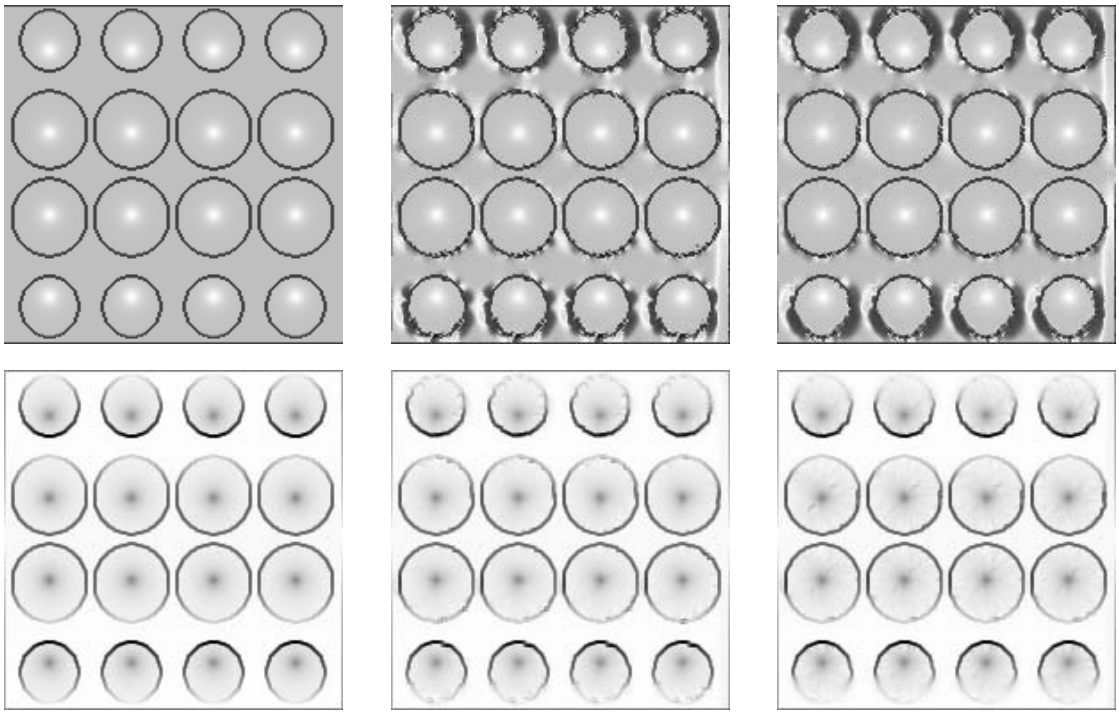

Figure 6.15: The shape-index (row 1) and the curvedness (row 2) images of the ground-truth (left) and those using the uncompensated (middle) and compensated (right) Hessian.

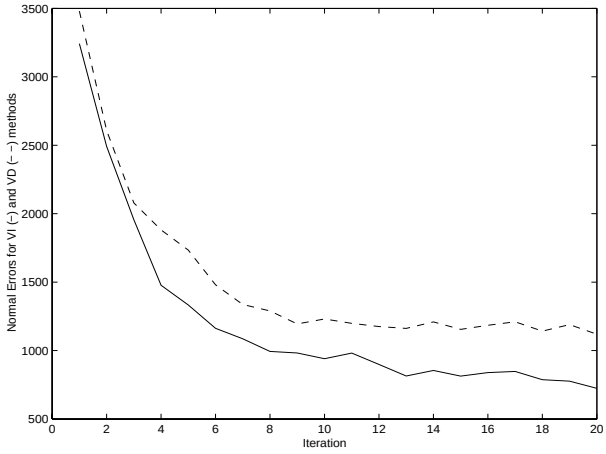

Figure 6.16: TMS surface normal error versus iteration number obtained using the uncompensated (dashed curve) and compensated (solid curve) Hessian methods.

iteration No. 20) in Figure 6.13 and the surface normals for other iterations are not shown here. The TMS error is found by summing the squared differences between every recovered surface normal and the corresponding ground-truth surface normal over the whole surface. Here, the solid curve is for the compensated Hessian whereas the dashed curve is for the uncompensated Hessian. It is clear from the curves that the compensated Hessian

results in smaller TMS errors compared to the uncompensated Hessian.

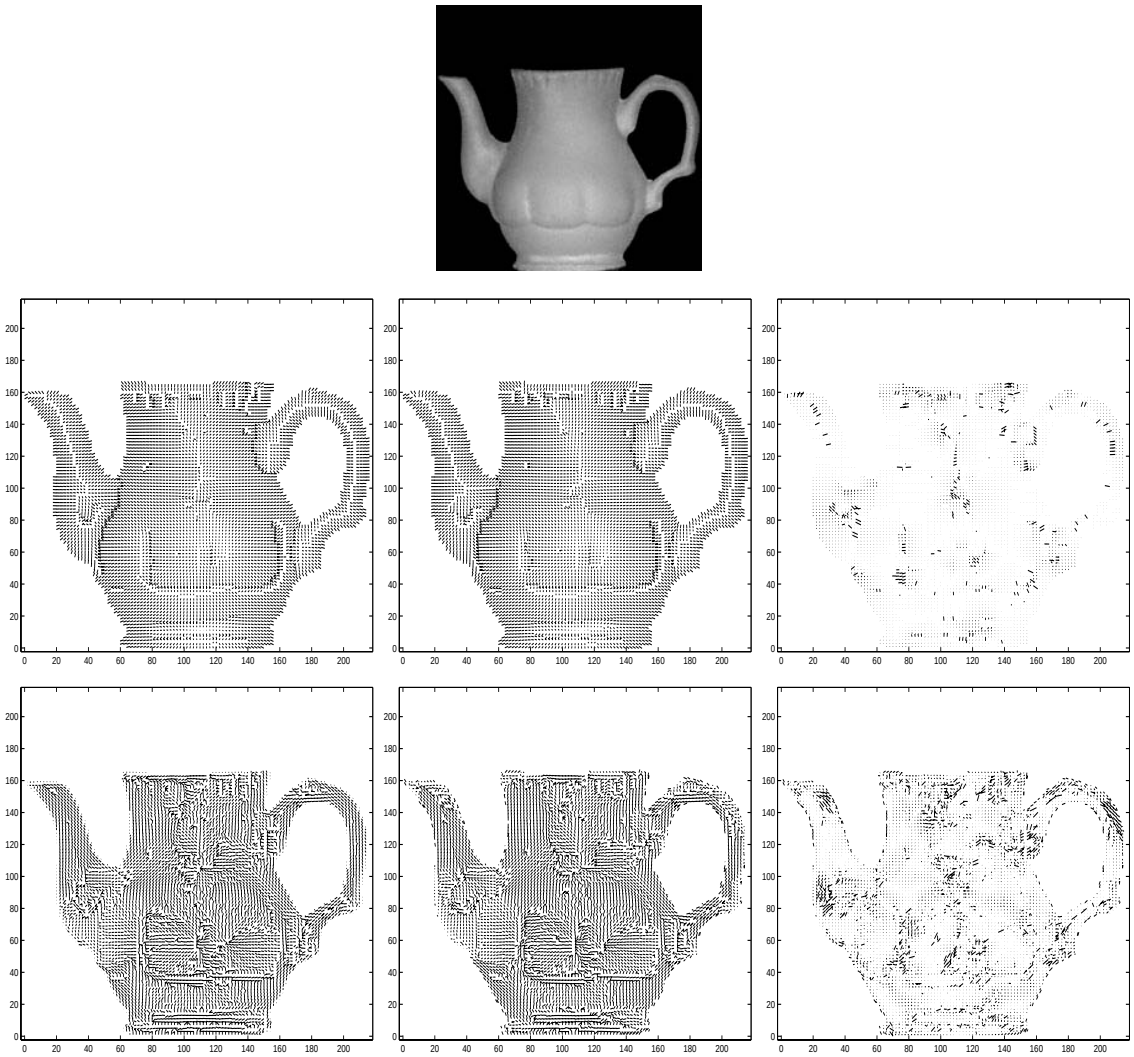

Figure 6.17: A terra-cotta coffee-pot: fields of surface normals (row 2) and minimum curvature directions (row 3) obtained using the uncompensated (left) and compensated (middle) Hessian in the D-S method; (right) differences between the two sets of vectors.

Finally, we experiment with a real-world image. In the top row of Figure 6.17, we show the brightness data. The fields of surface normals and principal curvature directions are shown in the second and the third rows. The first column is for the results obtained using the uncompensated Hessian, while the second column is for those obtained using the viewpoint-compensated Hessian. In the third column of the figure, we show the difference vector fields. From the difference fields, it is clear that differences in both surface normals and principal curvature directions are greatest in the proximity of highly curved regions

of the object surface. It is here that the rate of change of surface inclination is greatest.

6.6 Conclusions

To conclude this chapter, we have presented a new smoothing method for use in conjunction with SFS. The method adjusts the surface normal direction in a manner which is consistent with the principal curvature directions. Since smoothing of the principal curvature field is not appropriate at saddles or umbilics, we control the process using a topography sensitive process. The net effect is a smoothing process which allows the principal curvature directions to take control in parabolic surface regions. The net result is a SFS method which improves effectiveness in detecting crest lines.

Furthermore, we have shown how the viewpoint compensation method of Woodham can be used to compute improved estimates of the Hessian from the Gauss map in surface in SFS. This improved estimate has been used for the purposes of curvature sensitive smoothing. We demonstrate, that the compensated Hessian overcomes problems associated with overestimating curvature at locations of significant surface inclination.

There are clearly a number of ways in which the work described in this chapter can be extended. In the first instance, the smoothing process can be extended to elliptic and hyperbolic surface regions using curl and divergence criteria to control the regularization of the principal curvature directions. A second line of investigation will be to incorporate the smoothing process into situations where the reflectance is not completely Lambertian and different geometric constraints need to be applied to the surface normals.

Chapter 7

Conclusions

In this chapter, our first aim is to assess the strengths and weaknesses of the contributions in this thesis. Our second aim is to discuss critically some possible ways in which the weaknesses may be overcome. Finally, we conclude how the work presented in the thesis can be used as the basis for further research.

7.1 Strengths and Weaknesses of the Thesis

The thesis makes four contributions. The first contribution is that we have developed a shape-from-shading (SFS) method which is aimed at recovering surface normals for surfaces with specular highlights. The reflectance models used in this method are the Lambertian diffuse reflectance model and the Torrance-Sparrow specular reflectance model. To combine the Lambertian and specular surface normals at each iteration, we have adopted a maximum *a posteriori* probability estimation framework using the iterative conditional modes algorithm of Besag. We have shown that the method improves the accuracy of the recovered surface normals when compared with a recent SFS method which cannot deal with specularities. The final surface normals obtained by the method allow us to reconstruct the specular reflectance image using the Torrance-Sparrow model. Finally, using the reconstructed specular image, we have presented a method for removing specular highlights from the raw image and obtaining the matte image. We have tested three

different distribution functions to model the microfacet slopes. Depending on the value of its parameter, each function produces specific patterns for the *a posteriori* probabilities. However, the final fields of posterior-mean surface normals are quite similar. Hence, although these functions result in slightly different reconstructed specular intensity, the separated matte images are almost identical.

The second contribution is that we have proposed a method for recovering the Lambertian image from a non-Lambertian matte image. The author believes that the work is important because the bulk of the recent work in SFS restricts itself to a Lambertian model due to its simple form and that it only depends on incidence angle. Hence, our aim in this method is to process the input matte image and obtain the corresponding Lambertian image as the compatible input to Lambertian SFS algorithms. We have used the Oren-Nayar model for rough surfaces, the Wolff model for smooth surfaces, and the combined model of Wolff-Nayar-Oren for surfaces of intermediate roughness. The process uses a lookup table approach rather than an analytic approach to find the cosine of the incidence angle from the corresponding brightness value at each point of the input image. For matte surfaces, the diffuse image is the same as raw image, while for shiny surfaces the diffuse image is obtained by removing specularities from the raw image using our highlight removal approach mentioned earlier. This photometric correction process can also be interpreted as recovering the incidence angles at every point of the surface from the non-Lambertian image data. These incidence angles are equivalent to the surface normal zenith angles. Hence, from this point of view, the process is actually part of the shape recovery process which recovers the surface normal zenith angles and leaves the recovery of the surface normal azimuth angles to the Lambertian SFS algorithm.

On the other hand, to compute a lookup table each model needs to be simplified so that the model can be expressed only in terms of incidence angle. The main assumption made here to achieve such expressions is that the light-source and viewing directions are almost identical. Under these assumptions the incidence and reflectance angles are almost identical, and so, the reflectance equation is a function of incidence angle. Each model has a parameter which specifies the degree of departure from the Lambertian model for dif-

ferent surfaces. The Oren-Nayar model has a roughness parameter, the Wolff model has an index of refraction parameter, and the combined model has both of these parameters. Here, we have used approximate values of these parameters for similar surfaces obtained from tabulations in the literature. We have used ground-truth data to compare the recovered Lambertian images using the three models. We have also investigated the effect of the correction process on the surface normal directions recovered using SFS. From the results of our experiments on a variety of synthetic and real-world surfaces, we conclude that when the appropriate model (together with a well-approximated value of its parameter) is used for correcting the matte image of a surface, both the corrected Lambertian image and the corrected surface normals show lower differences with the corresponding ground-truth data.

The third contribution is that we have used physics-based models for light scattering from rough surfaces to estimate the roughness properties of dielectrics and metals. One advantage of the proposed techniques for roughness estimation is that all reflectance measurements are performed using a digital camera and under the visible light, whereas the previous methods use light-sources with long wavelengths together with advanced instruments for reflectance measurement which are not usually used in computer vision. For dielectrics, we have used the Vernold-Harvey modification to the Beckmann-Kirchhoff model to make parameter estimates for slightly-rough or very-rough surfaces. For slightly-rough surfaces, we first estimate the RMS roughness using the model for the specular direction. This allows us to estimate the correlation length using the model for diffuse reflectance. For very-rough surfaces, however, we can only estimate the surface slope, i.e. the ratio of the RMS roughness to the correlation length. This work is novel and deserves further research. Finally, for metallic surfaces, both the Bennett-Porteus and the Torrance formulations of the Davies theory can be used to estimate the RMS roughness and the correlation length. For metallic surfaces there are similar methods in the literature using the same models which work at long wavelengths.

These estimation techniques allow us to exploit the physics-based reflectance models for further surface analysis. For instance, for dielectric surfaces once the corresponding

estimates are found, the modified model of Beckmann-Kirchhoff can be used for three different applications. The two straightforward applications are surface re-illumination and rendering. Both of these applications use the general forms of the model variants. The third application is to use the fitted model to the reflectance data in a lookup table to recover the Lambertian image from the diffuse image of a curved surface with similar roughness properties. Again, we can perform this photometric correction process only under conditions of identical light-source and viewing directions.

The fourth, and final, contribution of the thesis is that we have introduced a Lambertian SFS method. Since, we have developed novel techniques for recovering the Lambertian component from the raw image, we make this correction more effective by developing a Lambertian SFS method in order to improve the process of shape recovery as well. Our method is an extension of the Worthington and Hancock algorithm (Worthington and Hancock, 1999). The surface normals are constrained to fall on a cone whose axis points in the light-source direction and whose apex angle is the inverse cosine of the Lambertian reflectance. However, rather than smoothing the field of surface normals using robust error kernels, we smooth the directions of the local Darboux frames using the field of principal curvature directions. Hence, we circumvent the problem of solving an under-constrained inverse reflectance equation to recover surface normal direction. It is this inverse problem that renders the non-Lambertian reflectance models intractable. By utilizing the field of principal curvature directions, we recover more accurate surface normals both at parabolic locations on the surface and at the locations of surface discontinuities (e.g. intersections between shapes).

7.2 Suggestions for Further Research

In this section we describe how the weaknesses of the work reported in the thesis can be overcome. Moreover, following the research presented in this thesis, some possible directions for future research are suggested.

7.2.1 Shape from Multiple Images using Physical Reflectance Models

One of the weaknesses of the thesis is the assumptions made in the photometric correction process in Chapters 4 and 5. The most limiting assumption is that the light-source direction is almost identical to the viewing direction. Although under these conditions the improved diffuse reflectance models can be used to recover the Lambertian reflectance from the diffuse reflectance, the method fails to do so for other lighting conditions.

Here we suggest a method to extend the photometric correction process to any lighting conditions. One reason why complex reflectance models are intractable for the purpose of shape recovery is that they depend on both incidence and reflectance angles. Specifically, when only one input image is used, the SFS problem is under-constrained. Hence, we suggest to use more than one image taken from an identical camera position but under different light-source directions. This approach is similar to photometric stereo and so allows complex reflectance models to be rendered tractable. If this approach is adopted then photometric correction is no longer a process separate from the SFS algorithm. A diffuse reflectance model, either physical or phenomenological, may be used in this method according to the scales of roughness of the surface. In what follows, we present a description of this proposed model-based shape from photometric stereo method.

The method commences by estimating an initial value for the incidence angle θ_i^0 at each point of the surface. Here, it is assumed that the relation between the initial set of surface normals and the image data is expressed by Lambert's law. This means that if the image intensity I_0 is used as the first input, at each image location we can write $\theta_i^0 = \cos^{-1}(I_0)$.

The next step is to find all possible surface normals at each single point using the non-Lambertian reflectance models. As described in Chapter 4, the effect of different reflectance models in the correction process is that the actual incidence angle θ_i deviates from the Lambertian incidence angle, i.e. θ_i^0. The difference in value between the non-Lambertian incidence angle and θ_i^0 depends on the model and the value of the roughness parameter used. However, there is definitely a maximum value for these difference values

which is denoted by $\Delta\theta_i^{max}$. Hence, the range of non-Lambertian incidence angles is $[\theta_i^0 - \Delta\theta_i^{max}, \theta_i^0 + \Delta\theta_i^{max}]$. If we divide $\Delta\theta_i^{max}$ into m equal intervals so that $\Delta\theta_i^{max} = m\Delta\theta_i^{min}$, then there will be $2m$ possible incidence angles corresponding to $2m$ possible surface normals. In other words, the two complementary incidence angles in this range are given by $\theta_i^{\pm k} = \theta_i^0 \pm k\Delta\theta_i^{min}$ where k is in the range $[1, m]$. Note that the number m depends on the required precision of the results.

So far, we have only found the incidence angles corresponding to the possible surface normals at each single point on the surface. Now it is time to compute an extended form of our lookup table using non-Lambertian diffuse reflectance models. This new lookup table has two inputs and three outputs. Its first input is the image intensity I_t and its second input is the incidence angle θ_t, whereas its three outputs are the three components of the corresponding surface normal, i.e. $< N_x, N_y, N_z >$. To produce this 3D lookup table, we use a half-sphere which has every possible surface normal in a 3-D coordinate system. Recall that a half-cylinder was used to produce the 2D lookup table. The viewing direction is a fixed direction aligned in the z axis direction, while the light-source direction can be varied and also measured using a constrained experimental setup. For every light-source direction, a lookup table can be computed for the reflectance model used. When computing the lookup table, we use all the surface normals on the half-sphere to compute the corresponding reflectance values using the model. Again, the number of surface normals used depends on the required precision of the results. With each lookup table to hand, we can find at most $2m$ possible surface normals at each image location.

The final step of the shape recovery algorithm is to use a selection, and/or rejection, method so that the most accurate set of surface normals can be obtained. If only one input image is available, the selection method may emphasize surface normal consistency and also image gradients to obtain the most likely surface normals in the proximity of small neighborhoods. However, another option is to use further input images with different light-source directions. In this case, for each image corresponding to an individual light-source, an individual lookup table should be computed. Hence, in addition to neighborhoods, the selection process can also focus on corresponding pixels in different images.

Note that each set of possible surface normals correspond to an identical deviation from the Lambertian incidence angles. However, it is not necessary that every selected surface normal belong to an identical set of surface normals. Note that this last feature provides extra flexibility to the algorithm and results in the recovery of more accurate shape. It is also interesting to investigate whether this algorithm can be modified to follow a similar approach to the parallel or cascade photometric stereo of Zhang et al. (Zhang et al., 1996; Zhang and Shah, 1997).

7.2.2 Roughness Estimation for the Oren-Nayar Model

The photometric correction process described in Chapter 4 can be improved by estimating the roughness parameter used in the simplified form of the Oren-Nayar model or in the combined model of Wolff-Nayar-Oren. Here we suggest two different methods for estimating this parameter.

The first method is based on fitting the model to the reflectance data. Here, we need to use a section of the object surface which is approximately cylindrical. Using this cylindrical section, we can measure the mean intensity values across the direction perpendicular to the cylinder axis. The data obtained in this way can be plotted versus the incidence angle values found from geometry. Once the curve corresponding to the real-word data is obtained, the closest possible curve obtained using the model can be fitted to it. The value of the roughness parameter used in the fitted model is the estimated value and can be used for photometric correction for the whole surface object of any shape.

The second method is to approximate the model given by Equation (4.9) for two different surface geometries. Specifically, for the locations on the surface where the incidence and reflectance angles are close to zero, the second term $B\sin(\theta_i)$ is very small and so we can write $I_A = I(\theta_i \approx 0) \approx A$. Also, for the limbs of the surface where the incidence and reflectance angles are close to $\pi/2$, it follows that $I_B = I(\theta_i \approx \pi/2) \approx B$. Hence, we can find these values for several locations on the surface and compute average values for the two quantities I_A and I_B. The roughness parameter can then be estimated using

the equation $I_B/I_A = B/A$ where the only unknown is the roughness parameter σ_a.

On the other hand, it is worth mentioning that recently some useful data including the parameter values of the Oren-Nayar model extracted from a variety of rough surface samples are provided in the "Columbia-Utrecht Reflectance and Texture Database" (Dana et al., 1996). Rather than estimating the parameters for surfaces with unknown parameter values, the values given in this database can be used together with surfaces similar to the corresponding surface samples.

Note that, we do not need to estimate the index of refraction parameter n which is used in the Wolff model. The first reason for this is that the value of n is limited in the range $[1.4, 2.0]$ and when either of the two limits 1.4 or 2.0 is used, the difference between the corresponding reflectance values is negligible. The second reason is that the value of n can be found in the literature just by referring to the surface material.

7.2.3 Testing Alternative Reflectance Models

As an alternative to the models studied in this thesis, it would be interesting to study the reflectance model of van Ginneken et al. (van Ginneken et al., 1998) for the purpose of photometric correction. Fortunately, since this model uses the distribution of the heights and the correlation function to derive the distribution of slopes, it is similar to the models from physical optics. On the other hand, the model includes the effects of masking and shadowing using geometrical optics. Hence, a comparison between this model and different variants of the modified B-K model would be interesting. For instance, the estimates of surface slope parameter used in this model can be compared with the estimates of surface slope parameter (σ/T) used in the B-K model. However, it is not a straightforward task to implement the model since the diffuse radiance component contains an integral which cannot be evaluated analytically.

Similar to most of the other reflectance models which are equated for total radiance, the second problem with the model of van Ginneken et al. is the problem of balancing between the specular and diffuse reflectance components. This problem may be overcome

using an alternative reflectance model which provides a unique equation for the two reflectance components. One such model has been recently developed by Sun (Sun, 2003) using physical parameters. The model is applicable to both smooth and rough surfaces. Specifically, the model generates a sharp specular highlight for smooth surfaces and gives a diffuse behavior for rough surfaces. This model also has an advantage over the B-K model since it provides a single solution for all surface roughness conditions.

There are also alternative purely physics-based reflectance models that may be tested. For instance, Bahar and Chakrabarti (Bahar and Chakrabarti, 1987) have developed a reflectance model for the graphics community based on a full-wave analysis of electromagnetic scattering by rough surfaces. Their model accounts for both specular and diffuse reflection and is relatively easy to implement. It would be interesting to compare this model with the original and modified B-K model.

7.2.4 Modifying the Beckmann Model for Missing Effects

Although the modification of Vernold-Harvey to the B-K model delivers more accurate reflectance predictions than the original B-K model, it can be criticized in two counts. Firstly, the modification is not completely based on physical optics theory. Instead, the modification relies on empirical and phenomenological arguments. This can be considered as a step backwards specifically, since they suggest the replacement of the geometrical factor, which depends on both incidence and reflectance angles, with a simplistic Lambertian cosine term. Secondly, while the modified model is aimed at compensating for multiple scattering and shadowing effects which are not included in the original model, it still fails to fit to real-world reflectance data at incidence angles corresponding to occluding boundaries of objects.

On the other hand, the B-K model is limited to very-rough or slightly-rough surfaces and it suffers from the lack of a specific variant for moderately-rough surfaces. To address this omission, we (Ragheb and Hancock, 2004) have recently reported the results of a study aimed at modifying the B-K model using the Fresnel coefficient and suggesting

a variant for medium roughness. We believe that there is still great potential for further research towards finding alternative models to replace the geometrical factor in order to improve the predictions of the B-K reflectance model. Obviously one fundamental modification would be to start from the Helmholtz integral equation and extend Beckmann's method by considering further modifications. There has been reasonable interest in this problem. For instance, Bruce and Dainty (Bruce and Dainty, 1991) refined the B-K theory by including shadow functions and Fresnel coefficients for both dielectrics and metals. They have also extended the theory to account for multiple scattering effects.

7.2.5 Reflectance Data Against Reflectance Models

An important step in bringing physics-based reflectance models, and specifically the B-K model, to the computer vision community is to show that they are applicable to real-world images of different surface materials and shapes. However, extracting reflectance data against incidence, and/or reflectance, angle for curved objects using just a single image is not an straightforward task. To perform a comparison between the measured reflectance data and reflectance models, we have used cylindrical surfaces so that we could approximate an incidence angle corresponding to each image brightness value. However, the standard way of doing such comparison is to measure BRDF's. Although it is difficult to measure, the BRDF has been recently and extensively studied in computer vision. For instance, Dana et al. (Dana et al., 1997) have catalogued the BRDF's of several types of 3D surface textures. It is well known that surfaces with large scales of roughness are mainly textured in appearance. Hence, we suggest a comparison between some recent BRDF measurements and physics-based reflectance models for very-rough surfaces. One example for this work can be found in the recent paper by Watson and Raven (Watson and Raven, 2001).

There are alternative approaches to BRDF measurements. For instance, Robles-Kelly and Hancock (Robles-Kelly and Hancock, 2003) have recently reported a method that can be used to estimate the surface radiance function from complex surfaces without the

need to accurately measure the four angular variables, and without the need for accurate light-source or camera calibration. The method can be applied to surfaces with isotropic and homogeneous reflectance properties. The method applies when the viewer and light-source directions are identical, and makes use of the cumulative distribution of image intensity gradients to estimate the surface radiance function. The surface radiance estimates obtained using this method can be used to test various reflectance models described in the thesis including the B-K model. We (Ragheb et al., 2004) have recently reported some results comparing the empirical radiance data from this method with the predictions delivered by existing reflectance models for smooth and rough surfaces.

Finally, it is important to stress that since very-rough surfaces exhibit 3D texture appearance, the work presented in Chapters 4 and 5 may be viewed as a promising step to bridge the gap between shape-from-shading and shape-from-texture. However, this topic is beyond the scope of the thesis and deserves further research.

Bibliography

Angelopoulou, E. and Poger, S. (2003). The color of specular highlights. In *Proc. of the SPIE: Human Vision and Electronic Imaging VIII*, volume 5007, pages 298–309.

Bahar, E. and Chakrabarti, S. (1987). Full-wave theory applied to computer-aided graphics for 3d objects. *IEEE Computer Graphics and Applications*, 7(7):46–60.

Barrick, D. E. (1968). Rough surfaces scattering based on the specular point theory. *IEEE Trans. on Antennas Propagation*, 16:449–454.

Beckmann, P. (1967). *Scattering of Light by Rough Surfaces: Progress in Optics, Volume VI, Edition E Wolf*, pages 55–69. Elsevier Science Publishers, Amsterdam.

Beckmann, P. and Spizzichino, A. (1963). *The Scattering of Electromagnetic Waves from Rough Surfaces*. Pergamon Press Ltd., New York.

Belhumeur, P. N. and Kriegman, D. J. (1996). What is the set of images of an object under all possible lighting conditions? In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 270–277.

Bennett, H. E. and Porteus, J. O. (1961). Relation between surface roughness and specular reflectance at normal incidence. *Journal of Optical Society of America*, 51(2):123–129.

Bennett, J. M. and Mattsson, L. (1999). *Introduction to Surface Roughness and Scattering*. Second Edition. Optical Society of America, Washington, D.C.

Berman, D. H. and Perkins, J. S. (1985). Exponential substitution for kirchhoff scattering from gaussian rough surfaces. *Journal of Acoustical Society of America*, 78:1024–1028.

Besag, J. (1986). On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society of London*, 48:259–302.

Bichsel, M. and Pentland, A. P. (1992). A simple algorithm for shape from shading. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 459–465.

Black, M. J. and Rangarajan, A. (1996). On the unification of line processes, outlier rejection, and robust statistics with applications in early vision. *International Journal of Computer Vision*, 19(1):57–91.

Blake, A. and Brelstaff, G. (1988). Geometry from specularities. In *Proc. IEEE International Conference on Computer Vision*, pages 394–403.

Blake, A. and Bulthoff, H. (1991). Shape from specularities: computation and psychophysics. *Philosophical Transactions of the Royal Society of London*, 331:237–252.

Blinn, J. F. (1977). Models of light reflection for computer synthesized pictures. *Computer Graphics*, 11(2):192–198.

Bors, A. G., Hancock, E. R., and Wilson, R. C. (2003). Terrain analysis using radar shape-from-shading. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(8):974–992.

Brelstaff, G. and Blake, A. (1988). Detecting specular reflection using lambertian constraints. In *Proc. IEEE International Conference on Computer Vision*, pages 297–302.

Brooks, M. J. and Chojnacki, W. (1994). Direct computation of shape from shading. In *Proc. International Conference on Pattern Recognition*, pages 114–119.

Brooks, M. J. and Horn, B. K. P. (1985). Shape and source from shading. In *Proc. International Joint Conference on Artificial Intelligence*, pages 932–936.

Bruce, N. C. and Dainty, J. C. (1991). Multiple scattering from rough dielectric and metal surfaces using the kirchhoff approximation. *Modern Optics*, 38(8):1471–1481.

Bruckstein, A. M. (1988). On shape from shading. *Computer Vision, Graphics and Image Processing*, 44:139–154.

Caron, J., Lafait, J., and Andraud, C. (2002). Scalar kirchhoff's model for light scattering from dielectric random rough surfaces. *Optics Communications*, 207:17–28.

Chandley, P. J. (1976). Surface roughness measurements from coherent light scattering. *Optics and Quantum Electron.*, 8:323–327.

Chantler, M. and McGunnigle, G. (2000a). On the use of gradient space eigenvalues for rotation invariant texture classification. In *Proc. International Conference on Pattern Recognition*, pages 943–946.

Chantler, M. and McGunnigle, G. (2000b). The response of texture features to illuminant rotation. In *Proc. International Conference on Pattern Recognition*, pages 955–958.

Chantler, M., Schmidt, M., Petrou, M., and McGunnigle, G. (2002). The effect of illuminant rotation on texture filters: Lissajous's ellipses. In *Proc. European Conference on Computer Vision*, pages 289–303.

Clark, J. J. and Yuille, A. L. (1990). Shape from shading via the fusion of specular and lambertian image components. In *Proc. International Conference on Pattern Recognition*, pages 88–92.

Cook, R. L. and Torrance, K. E. (1981). A reflectance model for computer graphics. *Computer Graphics*, 15(3):307–316.

Cryer, J. E., Tsai, P. S., and Shah, M. (1995). Integration of shape from shading and stereo. *Pattern Recognition*, 28(7):1033–1043.

Curran, W. and Johnston, A. (1996). The effect of illuminant position on perceived curvature. *Vision Research*, 36(10):1399–1410.

Dana, K. J. and Nayar, S. K. (1998). Histogram model for 3d textures. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 618–624.

Dana, K. J. and Nayar, S. K. (1999). Correlation model for 3d texture. In *Proc. IEEE International Conference on Computer Vision*, pages 1061–1067.

Dana, K. J., Nayar, S. K., van Ginneken, B., and Koenderink, J. J. (1996). Reflectance and texture of real-world surfaces. Technical report, Columbia University. No. CUCS-048-96 (www.cs.columbia.edu/CAVE/curet/).

Dana, K. J., Nayar, S. K., van Ginneken, B., and Koenderink, J. J. (1997). Reflectance and texture of real-world surfaces. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 151–157.

Davies, H. (1954). The reflection of electromagnetic waves from a rough surface. In *Proc. IEEE part IV*, number 101, pages 209–214.

de Carmo, M. P. (1976). *Differential Geometry of Curves and Surfaces*. Prentice-Hall, NJ.

Depew, C. A. and Weir, R. D. (1971). Surface roughness determination by the measurement of reflectance. *Applied Optics*, 10(4):969–970.

Drbohlav, O. and Sara, R. (2002). Specularities reduce ambiguity of uncalibrated photometric stereo. In *Proc. European Conference on Computer Vision*, pages 46–60.

Dror, R. O., Adelson, E. H., and Willsky, A. S. (2001). Estimating surface reflectance properties from images under unknown illumination. In *Proc. of the SPIE: Human Vision and Electronic Imaging VI*, volume 4299, pages 231–242.

Dupuis, P. and Oliensis, J. (1992). Direct method for reconstructing shape from shading. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 453–458.

Elson, J. M., Bennett, H. E., and Bennett, J. M. (1979). *Scattering from Optical Surfaces: Applied Optics and Optical Engineering, Volume VII*, pages 191–244. Academic Press, Inc., New York.

Ferrie, F. P. and Lagarde, J. (1990). Curvature consistency improves local shading analysis. In *Proc. International Conference on Pattern Recognition*, pages 70–76.

Fesenkov, V. (1962). *Physics and Astronomy of the Moon*. Academic Press, New York.

Fleming, R. W., Dror, R. O., and Adelson, E. H. (2001). How do humans determine reflectance properties under unknown illumination? In *Proc. IEEE Computer Society Workshop on Identifying Objects Across Variations in Lighting: Psychophysics and Computation*, pages 3–10.

Frankot, R. T. and Chellappa, R. (1987). A method for enforcing integrability in shape from shading algorithms. In *Proc. IEEE International Conference on Computer Vision*, pages 118–125.

Frankot, R. T. and Chellappa, R. (1988). A method for enforcing integrability in shape from shading algorithms. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 10(4):439–451.

Hapke, B. W. (1963). A theoretical photometric function for the lunar surface. *Geophysical Research*, 68(15):4571–4586.

Harvey, J. E., Vernold, C. L., Krywonos, A., and Thompson, P. L. (1999). Diffracted radiance: A fundamental quantity in a non-paraxial scalar diffraction theory. *Applied Optics*, 38(31):6469–6481.

Hayakawa, H. (1994). Photometric stereo under a light source with arbitrary motion. *Journal of Optical Society of America A*, 11(11):3079–3089.

He, X. D., Torrance, K. E., Sillion, F. X., and Greenberg, D. P. (1991). A comprehensive physical model for light reflection. *ACM Computer Graphics*, 25(4):175–186.

Healey, G. H. and Binford, T. O. (1987). Local shape from specularity. In *Proc. IEEE International Conference on Computer Vision*, pages 151–160.

Horn, B. K. P. (1970). *Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View*. PhD thesis, Department of Electrical Engineering, Massachusetts Institute of Technology.

Horn, B. K. P. (1981). Hill shading and the reflectance map. *Proceedings of the IEEE*, 69(1):14–47.

Horn, B. K. P. (1990). Height and gradient from shading. *International Journal of Computer Vision*, 5(1):37–75.

Horn, B. K. P. and Brooks, M. J. (1989). *Shape from Shading*. MIT Press, Cambridge MA.

Ikeuchi, K. (1981). Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 3(6):661–669.

Ikeuchi, K. and Horn, B. K. P. (1981). Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17(1-3):141–184.

Kimmel, R. and Bruckstein, A. M. (1992). Shape from shading via level-sets. Technical report, Israel Institute of Technology. CIS Report 9209.

Kimmel, R. and Bruckstein, A. M. (1995a). Global shape from shading. *Computer Vision and Image Understanding*, 62(3):360–369.

Kimmel, R. and Bruckstein, A. M. (1995b). Tracking level-sets by level-sets: A method for solving the shape from shading problem. *Computer Vision and Image Understanding*, 62(1):47–58.

Kimmel, R., Siddiqi, K., Kimia, B. B., and Bruckstein, A. M. (1995). Shape from shading: Level set propagation and viscosity solutions. *International Journal of Computer Vision*, 16(2):107–133.

Koenderink, J. J. and van Doorn, A. J. (1992). Surface shape and curvature scales. *Image and Vision Computing*, 10:557–565.

Koenderink, J. J., van Doorn, A. J., and Kappers, A. M. L. (1992). Surface perception in pictures. *Perception and Psychophysics*, 52(5):487–496.

Krakauer, K. J. (1971). *Computer Analysis of Visual Properties of Curved Objects*. PhD thesis, Massachusetts Institute of Technology, AI-TR.

Langer, M. S. and Bulthoff, H. (2001). Human perception of local shape from shading under variable lighting. In *Proc. IEEE Computer Society Workshop on Identifying Objects Across Variations in Lighting: Psychophysics and Computation*.

Leclerc, Y. G. and Bobick, A. F. (1991). The direct computation of height from shading. In *Proc. IEEE Conference on Computer Vision And Pattern Recognition*, pages 552–558.

Lee, C. H. and Rosenfeld, A. (1985). Improved methods of estimating shape from shading using the light source coordinate system. *Artificial Intelligence*, 26(2):125–143.

Lee, K. M. and Kuo, C. C. J. (1992). Shape reconstruction from photometric stereo. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 479–484.

Lee, K. M. and Kuo, C. C. J. (1993). Shape from shading with a linear triangular element surface model. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 15(8):815–822.

Lee, S. W. and Bajcsy, R. (1992). Detection of specularity using color and multiple views. In *Proc. European Conference on Computer Vision*, pages 99–114.

Leung, T. and Malik, J. (2001). Representing and recognizing the visual appearance of materials using three-dimensional textons. *International Journal of Computer Vision*, 43(1):29–44.

Lin, S. and Lee, S. W. (1999a). Estimation of diffuse and specular appearance. In *Proc. IEEE International Conference on Computer Vision*, pages 855–860.

Lin, S. and Lee, S. W. (1999b). A representation of specular appearance. In *Proc. IEEE International Conference on Computer Vision*, pages 849–854.

Lin, S. and Lee, S. W. (2000). An appearance representation for multiple reflection components. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 105–110.

Lin, S., Li, Y., Kang, S. B., Tong, X., and Shum, H. Y. (2002). Diffuse-specular separation and depth recovery from image sequences. In *Proc. European Conference on Computer Vision*, pages 210–224.

Lin, S. and Shum, H. Y. (2001). Separation of diffuse and specular reflection from color images. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 341–346.

Magda, S., Kriegman, D., Zickler, T., and Belhumeur, P. (2001). Beyond lambert: Reconstructing surfaces with arbitrary brdfs. In *Proc. IEEE International Conference on Computer Vision*, pages 391–399.

Malik, J. (1989). Recovering three-dimensional shape from a single image of curved objects. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 11(6):555–566.

Mardia, K. (1999). Directional statistics and shape analysis. *Journal of Applied Statistics*, 26:949–957.

Marx, E. and Vorburger, T. V. (1990). Direct and inverse problems for light scattered by rough surfaces. *Applied Optics*, 29(25):3613–3626.

McGunnigle, G. and Chantler, M. J. (2000). Rough surface classification using point statistics from photometric stereo. *Pattern Recognition Letters*, 21(6-7):593–604.

Minnaert, M. (1961). *Planets and Satellites: the Solar System*, volume 3. University of Chicago Press, IL.

Miyazaki, D., Kagesawa, M., and Ikeuchi, K. (2004). Transparent surface modelling from a pair of polarization images. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 26(1):73–82.

Nayar, S. K., Fang, X. S., and Boult, T. (1993). Removal of specularities using color and polarization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 583–590.

Nayar, S. K., Ikeuchi, K., and Kanade, T. (1990). Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Trans. on Robotics and Automation*, 6(4):418–431.

Nayar, S. K., Ikeuchi, K., and Kanade, T. (1991). Surface reflection: Physical and geometrical perspectives. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(7):611–634.

Nene, S. A., Nayar, S. K., and Murase, H. (1996). Columbia object image library (coil-20). Technical report, Columbia University. No. CUCS-005-96.

O'Donnell, K. A. and Mendez, E. R. (1987). Experimental study of scattering from characterized random surfaces. *Journal of Optical Society of America A*, 4(7):1194–1205.

Ogilvy, J. A. (1991). *Theory of Wave Scattering from Random Rough Surfaces*. Adam Hilger, Bristol.

Okatani, T. and Deguchi, K. (2000). Closed form solution of local shape from shading at critical points. *International Journal of Computer Vision*, 40(2):169–178.

Oliensis, J. (1991a). Shape from shading as a partially well-constrained problem. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 559–564.

Oliensis, J. (1991b). Uniqueness in shape from shading. *International Journal of Computer Vision*, 6(2):75–104.

Oliensis, J. and Dupuis, P. (1993). A global algorithm for shape from shading. In *Proc. IEEE International Conference on Computer Vision*, pages 692–701.

Oren, M. and Nayar, S. K. (1995). Generalization of the lambertian model and implications for machine vision. *International Journal of Computer Vision*, 14(3):227–251.

Pentland, A. P. (1982). Finding the illuminant direction. *Journal of Optical Society of America*, 72(4):448–455.

Pentland, A. P. (1984). Local shading analysis. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 6(2):170–187.

Pentland, A. P. (1988). Shape information from shading: A theory about human perception. In *Proc. IEEE International Conference on Computer Vision*, pages 404–413.

Phong, B. T. (1975). Illumination for computer generated pictures. *Communications ACM*, 18:311–317.

Porteus, J. O. (1963). Relation between the height distribution of a rough surface and the reflectance at normal incidence. *Journal of Optical Society of America*, 53(12):1394–1402.

Prados, E., Faugeras, O. D., and Rouy, E. (2002). Shape from shading and viscosity solutions. In *Proc. European Conference on Computer Vision*, pages 790–804.

Ragheb, H. and Hancock, E. R. (2004). Fresnel correction of the beckmann model. submitted paper.

Ragheb, H., Robles-Kelly, A., and Hancock, E. R. (2004). Radiance data against reflectance models. In *Proc. Computer Vision Winter Workshop*, pages 119–128.

Rindfleish, T. (1966). Photometric method for lunar topography. *Photogrammetric Engineering*, 32(2):262–277.

Robles-Kelly, A. and Hancock, E. R. (2003). Lambertian correction for rough and specular surfaces. In *Proc. Video, Vision and Graphics*, pages 213–220.

Samaras, D. and Metaxas, D. (1999). Coupled lighting direction and shape estimation from single images. In *Proc. IEEE International Conference on Computer Vision*, pages 868–874.

Sander, P. T. and Zucker, S. W. (1990). Inferring surface trace and differential structure from 3d images. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(9):833–854.

Sara, R. (1994). *Local Shading Analysis via Isophotes Properties*. PhD thesis, Johannes Kepler University Linz, Department of System Sciences.

Solomon, F. and Ikeuchi, K. (1996). Extracting the shape and roughness of specular lobe objects using four light photometric stereo. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 18(4):449–454.

Stam, J. (1999). Diffraction shaders. In *Computer Graphics, Proc. ACM SIGGRAPH*, pages 101–110.

Stogryn, A. (1967). Electromagnetic scattering from rough finitely conducting surfaces. *Radio Science*, 2(4):415–428.

Stover, J. C. (1995). *Optical Scattering: Measurement and Analysis*. SPIE Optical Engineering Press, Washington.

Sun, Y. L. (2003). Statistically-based reflection model for rough surfaces. In *Proc. of the SPIE: Computational Imaging*, volume 5016, pages 91–102.

Swaminathan, R., Kang, S. B., Szeliski, R., Criminisi, A., and Nayar, S. K. (2002). On the motion and appearance of specularities in image sequences. In *Proc. European Conference on Computer Vision*, pages 508–523.

Szeliski, R. (1991). Fast shape from shading. *Computer Vision, Graphics and Image Processing*, 53(2):129–153.

Tagare, H. D. and deFigueiredo, R. J. P. (1991). A theory of photometric stereo for a class of diffuse non-lambertian surfaces. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(2):133–151.

Torrance, K. E. (1964). Monochromatic directional distribution of reflected thermal radiation from roughened dielectric surfaces. Master's thesis, University of Minnesota.

Torrance, K. E. and Sparrow, E. M. (1967). Theory for off-specular reflection from roughened surfaces. *Journal of Optical Society of America*, 57(9):1105–1114.

Tsai, P. S. and Shah, M. (1994). Shape from shading using linear approximation. *Image and Vision Computing*, 12(8):487–498.

van Ginneken, B., Stavridi, M., and Koenderink, J. J. (1998). Diffuse and specular reflectance from rough surfaces. *Applied Optics*, 37(1):130–139.

Vega, O. E. and Yang, Y. H. (1993). Shading logic: A heuristic approach to recover shape from shading. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 15(6):592–597.

Vernold, C. L. and Harvey, J. E. (1998). A modified beckmann-kirchhoff scattering theory for non-paraxial angles. In *Proc. of the SPIE: Scattering and Surface Roughness II*, volume 3426, pages 51–56.

Watson, R. M. and Raven, P. N. (2001). Comparison of measured brdf data with parameterized reflectance models. In *Proc. of the SPIE: Targets and Backgrounds VII: Characterization and Representation*, volume 4370, pages 159–168.

Wei, G. Q. and Hirzinger, G. (1997). Parametric shape-from-shading by radial basis functions. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 19(4):353–365.

Wolff, L. B. (1994a). Diffuse-reflectance model for smooth dielectric surfaces. *Journal of Optical Society of America A*, 11(11):2956–2968.

Wolff, L. B. (1994b). On the relative brightness of specular and diffuse reflection. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 369–376.

Wolff, L. B., Nayar, S. K., and Oren, M. (1998). Improved diffuse reflection models for computer vision. *International Journal of Computer Vision*, 30(1):55–71.

Woodham, R. J. (1980). Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144.

Woodham, R. J. (1981). Analyzing images of curved surfaces. *Artificial Intelligence*, 17:117–140.

Woodham, R. J. (1994). Gradient and curvature from the photometric-stereo method, including local confidence estimation. *Journal of Optical Society of America A*, 11(11):3050–3068.

Worthington, P. L. and Hancock, E. R. (1999). New constraints on data-closeness and needle-map consistency for sfs. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(11):1250–1267.

Yuille, A. L. and Snow, D. (1997). Shape and albedo from multiple images using integrability. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 158–164.

Zhang, R. and Shah, M. (1997). Iterative shape recovery from multiple images. *Image and Vision Computing*, 15(11):801–814.

Zhang, R., Tsai, P., Cryer, J. E., and Shah, M. (1999). Shape from shading: A survey. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(8):690–706.

Zhang, R., Tsai, P. S., and Shah, M. (1996). Photomotion. *Computer Vision and Image Understanding*, 63(2):221–231.

Zheng, Q. and Chellappa, R. (1991). Estimation of illuminant direction, albedo, and shape from shading. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(7):680–702.

Zickler, T. E., Belhumeur, P. N., and Kriegman, D. J. (2002). Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *International Journal of Computer Vision*, 49(2-3):215–227.