

Imperial College of Science, Technology and Medicine
Department of Computing

On-site Surface Reflectometry

by

Jérémy Maxime Riviere

Submitted in part fulfilment of the requirements for the
degree of
Doctor of Philosophy in Computing
and the Diploma of Imperial College, 2017

Statement of originality: I, Jérémy Maxime Riviere, hereby declare that the work presented in this thesis is my own and that any third party material has been adequately referenced.

©Jérémy Riviere 2017

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

“La Terre est bleue comme une orange.” - Paul Éluard.

Abstract

The rapid development of Augmented Reality (AR) and Virtual Reality (VR) applications over the past years has created the need to quickly and accurately scan the real world to populate immersive, realistic virtual environments for the end user to enjoy. While geometry processing has already gone a long way towards that goal, with self-contained solutions commercially available for on-site acquisition of large scale 3D models, capturing the appearance of the materials that compose those models remains an open problem in general uncontrolled environments.

The appearance of a material is indeed a complex function of its geometry, intrinsic physical properties and furthermore depends on the illumination conditions in which it is observed, thus traditionally limiting the scope of reflectometry to highly controlled lighting conditions in a laboratory setup. With the rapid development of digital photography, especially on mobile devices, a new trend in the appearance modelling community has emerged, that investigates novel acquisition methods and algorithms to relax the hard constraints imposed by laboratory-like setups, for easy use by digital artists. While arguably not as accurate, we demonstrate the ability of such self-contained methods to enable quick and easy solutions for on-site reflectometry, able to produce compelling, photo-realistic imagery.

In particular, this dissertation investigates novel methods for on-site acquisition of surface reflectance based on off-the-shelf, commodity hardware. We successfully demonstrate how a mobile device can be utilised to capture high quality reflectance maps of spatially-varying planar surfaces in general indoor lighting conditions. We further present a novel methodology for the acquisition of highly detailed reflectance maps of permanent on-site, outdoor surfaces by exploiting polarisation from reflection under natural illumination.

We demonstrate the versatility of the presented approaches by scanning various surfaces from the real world and show good qualitative and quantitative agreement with existing methods for appearance acquisition employing controlled or semi-controlled illumination setups.

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my supervisor Dr Abhijeet Ghosh for introducing me to the exciting discipline that is appearance modelling. Throughout my years as a PhD student, I have learned a great deal from him on to conduct research and his expertise in appearance modelling has greatly helped shape this thesis. Abhijeet, thank you for being a great mentor and always pushing me to become a better researcher. I also take this opportunity to thank Pieter Peers and Ilya Reshetouski with whom I have collaborated on the projects that make up this thesis, and Luka Filipi for his great job on the mobile polarimetric application. A special thanks to my examiners, Dr. Will Smith and Prof. Andrew Davison for their helpful suggestions on improving this dissertation.

I would also like to thank all the members of the Realistic Graphics and Imaging group, past and present: Daljit Singh Dhillon, Christos Kampouris, Jaewon Kim, Yiming Lin, Ilya Reshetouski and Antoine Toisoul for their role in nurturing a great working environment. A special thanks also to all my fellow students with whom I've shared my struggles, joy, meals and more: Chin Pang (Clint) Ho (also for the invaluable discussions on algebra and optimisation), Lukas Rupprecht, Raoul Franky Guiazon (ULU tmtc), Dr Mathieu Hu (pronounced "who", how cool is that?!), Xi Chen (to whom I owe quite a few of the pounds I've packed over the past 4 years), Silvia Vinyes Mora, Jean Kossaifi. Special thanks also to Lauriane Thorner and Christophe Riviere for proof-reading my thesis.

A shout-out to the people at Disney Research Zrich for making my internship a great experience and in particular Derek Bradley and Paulo Gotardo for the many night caps. Special thanks also to Thabo Beeler for a well appreciated trip to the zoo when I really needed it!

The light probe images used for environmental rendering throughout this dissertation are courtesy of Paul Debevec <http://www.pauldebevec.com/Probes/>.

Finalemment, j'aimerais profiter de cette opportunit  pour remercier mes amis et ma famille pour leur soutien au cours de ces ann es qui ont  t  parsem es de hauts et de bas, et en particulier mes parents, Marie-No lle et Jean-Claude Riviere et mon fr re, Christophe Riviere, qui ont toujours crus en moi et m'ont support s en toute situation. Cette th se leur est d di e.

Contents

Abstract	4
Acknowledgements	5
List Of Abbreviations	10
List Of Symbols	10
List Of Tables	12
List Of Figures	12
I Prologue	25
1 Introduction	27
1.1 Motivation	27
1.2 Appearance modelling in graphics	29
1.3 Thesis overview	30
2 Background And Related Work	34
2.1 The Bidirectional Reflectance Distribution Function (BRDF)	35
2.1.1 Fresnel reflectance	37
2.1.2 Specular reflection models	39
2.1.3 Microfacet distribution models	42
2.1.4 Diffuse reflection models	45
2.2 Reflectometry	46

2.2.1	Laboratory setup	46
2.2.2	Accessible reflectometry	47
2.2.3	Uncontrolled environments	50
2.3	Polarisation in vision and graphics	51
2.3.1	Reflectance separation/estimation	52
2.3.2	Shape estimation	54
	Prologue - Summary	57
	II Mobile Surface Reflectometry [1]	58
	3 Mobile Surface Reflectometry - Free-form Acquisition	60
3.1	Setup	61
3.2	Calibration	63
3.2.1	Data registration	63
3.2.2	Back-scattering direction estimation	64
3.2.3	Radiometric calibration	66
3.3	Reflectance recovery	67
3.3.1	Surface normal estimation	67
3.3.2	Diffuse component estimation	69
3.3.3	Specular component estimation	69
3.3.4	Results	70
3.4	Surface detail augmentation	74
3.4.1	Algorithm	74
3.4.2	Results	76
3.5	Discussion and limitations	77
	4 Mobile Surface Reflectometry - LCD-based Acquisition	79
4.1	Data acquisition	79
4.2	Processing pipeline	81
4.2.1	Reflectance recovery	82

4.2.2	Results	84
4.3	Appearance transfer	87
4.3.1	Algorithm	88
4.3.2	Results	90
Mobile Surface Reflectometry - Summary		93
III Outdoors Reflectometry [2]		94
Outdoors Reflectometry - Prelude		96
5 Polarisation In Graphics And Vision		98
5.1	Background	98
5.1.1	Mueller calculus	99
5.1.2	Polarising filters	101
5.1.3	Reflectors	104
5.2	Polarisation imaging	105
5.3	Image formation model	108
5.3.1	Partially linearly polarised incident illumination	108
5.3.2	Unpolarised incident illumination	110
5.3.3	Discussion	111
6 Polarisation Imaging Reflectometry In-the-wild		115
6.1	Digital Single-Lens Reflex (DSLR) setup	116
6.1.1	Data acquisition	116
6.1.2	Calibration	119
6.1.3	Polarisation processing pipeline	120
6.2	Reflectance recovery	122
6.2.1	Surface normal estimation	122
6.2.2	Diffuse albedo estimation	123
6.2.3	Specular component estimation	124

6.3	Results and analysis	127
6.3.1	Results	127
6.3.2	Discussion and error analysis	131
6.4	Mobile polarimetric setup	139
6.4.1	Acquisition pipeline	140
6.4.2	Results	143
	Outdoors Reflectometry - Summary	146
	IV Epilogue	147
	7 Conclusion	149
	Bibliography	152
	Appendices	163
A	Additional renderings - Mobile Surface Reflectometry (part II) . . .	165
B	Additional renderings - Outdoors reflectometry (part III)	172

List Of Abbreviations

RMS Root Mean Squared	IMU Inertial Measurement Unit
BTF Bidirectional Texture Function	AR Augmented Reality
BRDF Bidirectional Reflectance Distribution Function	VR Virtual Reality
dBRDF distribution-based BRDF	API Application Programming Interface
SVBRDF Spatially-Varying BRDF	PTAM Parallel Tracking And Mapping
TRS Transmitted Radiance Sinusoid	RANSAC RAndom SAmple Consensus
DOLP Degree Of Linear Polarisation	NDF Normal Distribution Function
DSLR Digital Single-Lens Reflex	Texel Texture Element
HDR High Dynamic Range	LUT Look-Up Table
CCD Charge-Coupled Device	SfP Shape from Polarisation
LCD Liquid-Crystal Display	SVD Singular Value Decomposition
SfM Structure from Motion	PDF Probability Density Function
CNN Convolutional Neural Network	RGB Red-Green-Blue
DOP Degree of Polarisation	SIFT Scale-invariant Feature Transform
PC Personal Computer	MVS Multi-View Stereo
MP MegaPixel	AEB Auto-Exposure Bracketting
HD High Definition	SNR Signal-to-Noise Ratio

List Of Symbols

\mathbf{Z} Set of natural integers	ρ_d Diffuse albedo
$f_r(\vec{\omega}_i, \vec{\omega}_o)$ BRDF	\vec{n} Surface normal
$F(\theta)$ Fresnel reflectance	σ Specular roughness
$G(\vec{\omega}_i, \vec{\omega}_o, \vec{n})$ Shadowing-masking term	Ω^+ Hemisphere of directions surrounding a surface patch oriented by its surface normal \vec{n}
$D(\vec{\omega}_i, \vec{\omega}_o, \vec{n}, \sigma)$ Microfacet distribution term	Ω Spherical domain of integration of the rendering equation
θ_B Brewster angle	L_i Incoming radiance
$f_d(\vec{\omega}_i, \vec{\omega}_o)$ Diffuse BRDF	$L_o(\vec{\omega}_o)$ Outgoing radiance
$f_s(\vec{\omega}_i, \vec{\omega}_o)$ Specular BRDF	$E_i(\vec{\omega}_i)$ Irradiance
ρ_s Specular albedo	
$F(0^\circ)$ Reflectance at normal incidence	

\mathbf{R}_{\parallel}	Reflectance of p-polarised (parallel) light as predicted by Fresnel equations	dinates
\mathbf{R}_{\perp}	Reflectance of s-polarised (perpendicular) light as predicted by Fresnel equations	χ Ellipticity angle expressed in local coordinates
δ	Phase retardation between the orthogonally projected images of the \vec{E} -vector	$\vec{\omega}$ Unit direction in spherical coordinates ($\vec{\omega} = (\theta, \phi)$)
$\mathbf{x}_{i,o,r,t}$	Subscript indicates input (i), output (o), reflection (r) or transmission (t), where x can be any symbol listed thereafter	\mathbf{s} Stokes vector ($s = [s_0, s_1, s_2, s_3]^T$)
\mathcal{P}	Degree Of Linear Polarisation (DOLP)	\vec{E}_{\parallel} Unit vector parallel to the plane of incidence
ψ	Angle of polarisation expressed in local coordinates	\vec{E}_{\perp} Unit vector perpendicular to the plane of incidence, such that $[\vec{E}_{\perp}, \vec{E}_{\parallel}, \vec{n}]$ form a right-handed orthonormal basis
		η Index of refraction

List of Tables

6.1	Statistical variation in surface normals of “red book” under different lighting conditions (left column), compared to two measurement methods employing controlled illumination.	133
-----	--	-----

List of Figures

2.1	BRDF geometry: The BRDF is a 4D real-valued function defined over the upper hemisphere (Ω^+) oriented by the normal to the surface (\vec{n}).	35
2.2	Diffuse vs Specular reflection: Notice how the diffuse material looks identical as the camera moves while the specular material changes rapidly as the camera is tilted to the right.	36
2.3	Fresnel visualisation: Geometry (a) and equations at a dielectric-dielectric interface (b).	38
2.4	Gonioreflectometer: Schematic of a typical gonioreflectometer as proposed by Murray-Colemann and Smith [3].	47
2.5	Degree of Polarisation (DOP) \mathcal{P}: Plots of the degree of specular polarisation (blue) and degree of diffuse polarisation (red) for unpolarised incident illumination reflected at an air-glass interface. .	53
2.6	Transmitted Radiance Sinusoid: The observation of the radiance resulting from specular reflection through a rotating linear polariser has the form of a phase-shifted sinusoid, where the phase (ϕ) is related to the azimuth of the surface normal ($\phi_{\vec{n}}$) by $\pm\frac{\pi}{2}$. . .	55

I	Examples of surface reflectance recovered using Mobile Surface Reflectometry: (a) A spatially varying rough specular material acquired using our hand-held free-form acquisition technique (chapter 3). (b) Highly specular surface reflectance recovered using mobile LCD-based reflectometry (chapter 4), with enhanced mesostructure from close-up observations under natural lighting (section 3.4). (c) Surface reflectance of a large spatially-varying material sample recovered using appearance transfer from surface reflectance obtained using the LCD-based approach for a small reference patch (section 4.3).	58
3.1	Acquisition setup: 10" Fujitsu Stylistic M532 Android tablet with an 8 MegaPixel (MP) back-facing camera with co-located LED flash.	61
3.2	Mobile Surface Reflectometry - free-from acquisition: Flash-based surface reflectometry in a dimly lit office room.	62
3.3	Calibration geometry: Device-centred coordinate system.	65
3.4	Back camera response: Given the limited control on exposure of our device, we instead fit a gamma function ($\gamma = 3.2$) to the observed radiance of the X-Rite ColorChecker®'s grayscale gradient.	67
3.5	Anniversary greeting card: Example reflectance maps (bottom) recovered from our mobile free-form acquisition setup by decomposing the measured reflectance trace (a) to recover per-Texture Element (Texel) diffuse albedo and specular reflectance (b,c), normals (d) and roughness (e).	68
3.6	Free-form acquisition - maps: Reflectance maps recovered from our free-form acquisition setup for three spatially-varying rough specular greeting cards.	71
3.7	Free-form acquisition - results: Photo-rendering comparison under point light illumination at normal incidence (back-scattered reflection).	73

3.8	Surface details enhancement input: We take two additional close-up observations of the sample under natural illumination through a window to recover additional mesoscopic surface details not visible in the reflectance maps recovered from free-form acquisition.	74
3.9	Surface detail enhancement algorithm: We add mesoscale details (middle column) to the maps recovered from free-form acquisition (first column) to obtain highly detailed reflectance maps (third column). Zoomed-in crop to highlight details.	75
3.10	Surface detail enhancement results: Here we compare photos of our samples (a) to renderings before (b) and after (c) surface detail enhancement. Zoomed in crop to highlight details.	76
3.11	Mobile Surface Reflectometry - validation: Comparison of renderings with normal maps obtained from sensor-based tracking (c) and 3D tracking (d), to a photograph (b). Here, we also demonstrate good photo-rendering matching for novel view renderings ((b)-(f)) in the mirror direction (a).	78
4.1	Mobile Liquid-Crystal Display (LCD)-based acquisition: We project gradient illumination patterns (b) from the device's screen and image the sample from the front facing camera (a).	80
4.2	Mobile LCD-based data: We acquire two sets of data under gradient illumination, taking advantage of the inherent polarisation of the device's screen for diffuse-specular separation.	81
4.3	Hybrid normals problem: Due to the low intensity of light emitted from the device's LCD screen and limited extent of the screen over the sample's upper hemisphere (Ω^+), the diffuse normals (a) are too noisy for rendering. Instead, we render both diffuse and specular reflections with specular normals (b).	83

- 4.4 **LCD-based mobile reflectometry - direct capture:** Reflectance maps acquired for small shiny objects as per the protocol outlined in section 4.1. The red line delimits the frontier between reflectance maps without post-processing (left) and reflectance maps obtained after surface detail augmentation as presented in section 3.4 (right). 84
- 4.5 **Bronze medal:** Reflectance maps ((a)-(d)) recovered from our mobile LCD-based acquisition setup. The maps allow good qualitative renderings (f) with highlights matching real photographs (e). 86
- 4.6 **Intel coin:** Reflectance maps ((a)-(d)) recovered from our mobile LCD-based acquisition setup. The maps allow good qualitative renderings (f) with highlights matching real photographs (e). 86
- 4.7 **Limited measurable size:** The non ideal overlap between view and light frustum (a) makes it impossible to directly measure larger samples (b). Our appearance transfer approach is directly motivated by the observation that spatially-varying specular appearance tends to be low-rank (c). 87
- 4.8 **Appearance transfer:** From a small crop obtained by direct capture ((a),(b)) and a specular cue texture (c) we transfer both roughness (d) and specular reflectance (e) properties to a larger portion of the sample. 88
- 4.9 **Appearance transfer - maps:** From observations of a small representative patch of the material (e-h), we transfer specular reflectance (c) and roughness (d) properties to the entire sample as per the protocol described in section 4.3.1. 91
- 4.10 **Appearance transfer - results:** Despite the simplicity of the method, our recovered reflectance maps produce compelling renderings (b,c) that well match real photographs (a). 92

II	Polarisation imaging reflectometry in-the-wild: High-resolution renderings from reflectance and normal maps of planar samples obtained by exploiting multiple polarisation observations under uncontrolled outdoor illumination.	94
5.1	Visualising polarisation: When looking down the propagation direction of a light wave, the tip of its electric field traces an ellipse as it oscillates (a). Stokes parameters are related to the parameters of the polarisation ellipse as per eq. (5.1). They span a 3D space represented in spherical coordinate on the Poincaré sphere (b). . . .	100
5.2	Linear polariser: Geometry of a general linear polariser rotated at an angle ϕ_o from its local coordinate system.	102
5.3	Circular polariser: Geometry of general circular polarisers.	103
5.4	Reflector: Geometry of a general reflecting optical element rotated by an angle $\phi_{\vec{n}}$ from its local coordinate system.	105
5.5	Polarisation imaging geometry: The angle of polarisation ψ_i of the incident beam of light is relative to the local coordinate system spanned by $[\vec{E}_{\perp}, \vec{E}_{\parallel}, \vec{\omega}_i]$. The outgoing coordinate system is defined by $[\vec{x}, \vec{y}, \vec{\omega}_o]$, the local coordinate system of the camera. ϕ_o is the angle of rotation of the linear polariser in front of the camera and ϕ is the angle between the camera's \vec{x} -axis and the direction perpendicular to the plane of incidence (\vec{E}_{\perp}).	108
6.1	Principal polarisation imaging setup: We employ commodity photography equipment often used for image-based lighting applications [4].	116
6.2	Polarisation imaging reflectometry in the wild: Measurement protocol. We acquire a set of High Dynamic Range (HDR) sequences near normal incidence and close to Brewster angle of incidence.	118

6.3	Semi-automatic registration tool: The user selects four matching corners on each frame and our tool automatically generates a regular grid of matched points ready to pass as input to VisualSfM.	119
6.4	Transmitted Radiance Sinusoid (TRS) fitting: For each near-Brewster view, we compute a per-pixel fit of eq. (5.22) to the acquired data.	120
6.5	Polarisation imaging reflectometry: Example reflectance maps recovered for a permanent on-site specular “drain cover”, captured on a busy side-walk close to Imperial College’s campus.	122
6.6	Change in reflected radiance $s_{r,0}$ due to incident polarised illumination ($DOP = 80\%$) is very similar for stainless steel (solid plots) and a dielectric ($\eta = 1.5$, dotted plots) around Brewster angle of incidence. The colours indicate three different angles of polarisation w.r.t. the plane of incidence (Red: $\psi_i = 0^\circ$, Green: $\psi_i = 45^\circ$, Blue: $\psi_i = 90^\circ$).	124
6.7	Reflectance maps ((a)-(d)) estimated from two views of the sample close to Brewster angle of incidence, under natural outdoors illumination (e). Our method is agnostic to the incident illumination and robust to changes in illumination during capture.	129
6.8	Comparisons of sample photographs (a) to matching renderings under the same incident illumination (b), as well as renderings in novel lighting environment (c).	130

- 6.9 **Brewster angle measurement validation:** Simulated TRS for a glass material (Index of Refraction $\eta = 1.5$) oriented at an azimuth $\phi_{\vec{n}} = 90^\circ$. First row: Simulation under unpolarised incident illumination - the maximum of the TRS is found at $\phi_o = 0^\circ$, as expected, for any angle of incidence θ_i . Second row: Simulation under partially linearly polarised illumination with a DOP of 80%. The different colours represent different angles of polarisation ψ_i . Unlike under unpolarised incident illumination, the phase of the TRS depends on the angle of polarisation of the incident illumination. However, behaviour similar to that obtained under unpolarised illumination can be observed again at and around Brewster angle (i.e when $\theta_i = \theta_B$). 132
- 6.10 **Diffuse-specular separation at normal incidence:** At 6pm, the sky is strongly linearly polarised at the zenith (a), which allows for good diffuse-specular separation ((c),(d)) near normal incidence. 134
- 6.11 **Reflectance maps comparison:** We compare the reflectance maps of the “red book” sample recovered with our method (first row) to those recovered from two methods under controlled illumination (second and third row). Our method shows good qualitative agreement with those methods. 135
- 6.12 Our method has problems with accurate estimation of surface reflectance over the metallic surface of this ornate book cover. . . . 137
- 6.13 **Partial result:** Bas-relief on an interior wall of an enclosed entrance to a courtyard acquired with our method. 139
- 6.14 **Mobile polarimetric setup:** We statically mounted a linear polariser to a mobile phone’s back camera and developed a custom application to allow mobile polarimetric measurements. The polariser’s optical axis is oriented parallel to the device’s landscape position. 140

6.15	Mobile polarimetric acquisition: The user rotates the device at 0° (a), 45° (b) and 90° (c) in order to sample the TRS. Notice how the intensity of the reflection on the “drain cover” changes with the orientation of the device.	141
6.16	Reflectance maps ((a)-(c)) estimated using our mobile polarimetric setup. As we cannot estimate the specular roughness, we borrowed it from our DSLR setup for the “drain cover” and manually set it to 0.2 for the “garden pavement” and sketch book for rendering (d).	143
6.17	Comparison of the reflectance maps obtained for the “drain cover” sample with our mobile polarimetric setup (top row) and DSLR setup (bottom row).	145
6.18	Comparison of the reflectance maps obtained for the “sketch book” sample with our mobile polarimetric setup and the LCD-based method of Ghosh et al. [5].	145
1	“Anniversary” greeting card: Environmental illumination renderings.	165
2	“New job” greeting card: Environmental illumination renderings.	166
3	“Twenty-one” greeting card: Environmental illumination renderings.	167
4	“Bronze medal”: Environmental illumination renderings.	168
5	“Cha” tea box cover: Environmental illumination renderings. . .	169
6	“Child’s book” cover: Environmental illumination renderings. . .	170
7	“Intel” coin: Environmental illumination renderings.	171
8	“Drain cover”: Environmental illumination renderings.	172
9	“Red book” cover: Environmental illumination renderings. . . .	173
10	“Red bricks”: Environmental illumination renderings.	174
11	“Garden pavement”: Environmental illumination renderings. . .	175
12	“Stone pavement”: Environmental illumination renderings. . . .	176
13	“Bas-relief”: Environmental illumination renderings.	177

14 **“Sketch book” cover:** Environmental illumination renderings. . . 178

PART I:

PROLOGUE

Chapter 1

Introduction

1.1 Motivation

Since its debut in the 1950s, the field of computer graphics has grown significantly from a relatively small discipline known only to a handful of researchers, to a core component of computer science, now pervasive in our society of mass media production. From its early age, a fundamental goal of computer graphics has been to produce photo-realistic imagery; that is, imagery that cannot be distinguished from real photographs. This immediately sparked interest among the movie and video game industries where the emphasis is on creating realistic, immersive contents for the end user. As the discipline developed, and the graphics technologies gradually became pervasive, a wider range of applications started to appear such as cultural heritage conservation, advertising, computer-aided analysis and recognition of real scenes, etc.

While a wide variety of models and algorithms exist for photo-realistic rendering, the end product of most physically-based engines greatly depends on the quality of input provided to the rendering software. The way we perceive an object is indeed a complex function of its geometry, intrinsic reflectance properties and illumination condition in which we observe it. Photo-realistic rendering thus requires expertise in multiple areas of science, such as computer vision and optics in order to recover high-resolution geometry, illumination and appearance which can

be challenging to express analytically. Researchers have thus started to investigate measurement-based methods for digitising their environment. In particular, in the remainder of this dissertation, we will be interested in appearance modelling and will investigate novel approaches for in-situ, image-based acquisition of surface reflectance using commodity hardware. This is in line with a recent trend in graphics to step away from the conventional controlled measurements requiring expensive laboratory setups and expert knowledge to drive such setups. The main goal in this novel trend is to allow quick and easy methods for reflectometry, where the emphasis is on producing compelling rendered imagery, often trading off accuracy in the strict radiometric sense.

Such methods based on commodity hardware have the potential to replace the otherwise cumbersome task incumbent upon digital artists to manually paint reflectance maps from scratch, for applications such as visual effects, virtual and augmented reality, etc. The added benefit of working with measured data lies in the ability to capture subtle details (scratches, rust, mould) to which the human perceptual system are particularly sensitive. Those subtleties are paramount to giving an impression of realism and can be otherwise complex to express analytically, procedurally and/or hand paint.

In general, recovering the appearance of an object is a complex, often ill-posed problem as it is a complex function of its geometry, reflectance properties and illumination conditions during observation. In our work, in order to make the problem tractable, we thus restrict ourselves to planar surfaces, allowing spatial variations in geometry that can be explained by a normal map: a texture that maps the local orientation of each surface point ($\vec{n} = \begin{bmatrix} x, y, z \end{bmatrix}$) to a Red-Green-Blue (RGB) colour¹. Furthermore, we restrict our discussion to surfaces exhibiting spatially-varying isotropic reflectance properties, well represented by the dichromatic reflectance model first proposed by Shafer in 1985 [6]. This model describes reflection as the sum of two independent components: a view-independent com-

¹Note that for display purposes, our normals are mapped to RGB triplets by the following formula: $\frac{\vec{n} + 1}{2} \rightarrow RGB$ which maps the up-vector $\vec{z} = [0, 0, 1]$ to the RGB colour $[0.5, 0.5, 1]$ and explains the overall pink-blue hue of all our normal maps

ponent that accounts for diffuse scattering of light by pigments of the material and a view-dependent component that accounts for specular reflection at the surface of the material. We model these effects with three additional textures: a diffuse albedo texture which encodes the matte aspect of the material and specular reflectance and roughness textures which encode the shininess of the material. Section 2.1 will discuss these aspects in more details.

1.2 Appearance modelling in graphics

The main goal in appearance modelling in general and in particular reflectometry is to measure and quantify a surface's appearance as a function of incident and outgoing directions. One of the most generic of these functions, the Bidirectional Texture Function (BTF), allows to capture realistic effects such as self-shadowing, self-occlusions and inter-reflections but often comes at a high computational cost and memory footprint. Instead in this dissertation, we will be interested in the Bidirectional Reflectance Distribution Function (BRDF) [7], a 4D function of incident and outgoing directions, which quantifies the complex interaction of light at the surface of a material in a more compact way (section 2.1). One of the early designs for BRDF measurement, the gonireflectometer [3, 8], consists of a point light source and reflectance detector, each mounted on a mechanical gantry, allowing for an exhaustive sampling of the hemisphere of incoming and outgoing directions around the surface. While yielding extremely accurate BRDF measurements, this design suffers from lengthy acquisition times as well as large memory footprint.

Over the past two decades, thanks to the advances in digital photography, the original design for the gonireflectometer has been widely revised and improved [9, 10], by taking advantage of image-based acquisition to speed-up the capture process. Another contributing factor in the development of image-based techniques for reflectance acquisition has been the seminal work of Debevec and Malik [11] on HDR imaging.

Furthermore, digital photography equipment has become pervasive, due largely to the developments in sensor technology allowing for an ever increasing quality in imagery at a lower cost. From high-end DSLR camera models released every other year to high-resolution cameras on mobile phones and tablets, high quality digital photography is nowadays accessible to anyone for virtually any budget. The availability of such technology has thus given rise to image-based appearance modelling techniques that exploit measurements from the real world to drive physically-based rendering engines.

In particular, a recent trend in the graphics community has been to develop simple and portable designs for reflectometry based on off-the-shelf commodity hardware [12, 13, 14], to enable non expert users to quickly and easily capture the appearance of real world material, simplifying the otherwise lengthy process required by digital artists, of hand painting reflectance maps. The work presented in this dissertation aims to further investigate methods for on-site reflectometry based on off-the-shelf commodity hardware, for use by non experts in the field of appearance modelling.

1.3 Thesis overview

The rest of this dissertation is organised as follows. First, we present some related work on appearance modelling and reflectometry by introducing the concept of the Bidirectional Reflectance Distribution Function (BRDF) (section 2.1). We then review some prior art on BRDF measurement (section 2.2) chronologically, from the very first designs to modern approaches based on commodity hardware, which our proposed mobile reflectometry approaches (part II) extend. We then end our related work section by reviewing previous work on polarisation in graphics and vision, which is a core component of our outdoors, passive reflectometry approach (part III).

Part II presents two novel methods for acquiring detailed spatially varying isotropic surface reflectance and mesostructure of a planar material sample using

commodity mobile devices. We first present a free-form hand-held method for the acquisition of reflectance maps of rough specular, spatially-varying planar surfaces, exploiting back-scattered measurements (chapter 3). Our second approach, targeted at highly specular materials, uses the LCD monitor as an extended source of illumination to illuminate the sample with polarised gradient illumination patterns (chapter 4), exploiting the inherent polarisation of the device’s screen for diffuse-specular separation. To overcome some of the limitations inherent to mobile acquisition in terms of spatial resolution, we further propose practical methods for appearance augmentation (section 3.4) and appearance transfer (section 4.3). The latter allows to hallucinate reflectance properties for samples too large to be captured directly by our LCD-based approach, from observations of a small representative crop of the sample.

The methods presented in part II are restricted to indoor environments where the room’s illumination can be controlled such that the illumination from the device dominates the ambient illumination, thus limiting their use outdoors where the ambient illumination is likely to be dominant, especially on a sunny day. It is therefore not possible to measure outdoor, on-site structures such as brick walls, pavements, etc. This limitation is common with most existing methods for reflectometry. To address this problem, we therefore present a novel approach for the acquisition of reflectance properties for permanent outdoors on-site planar surfaces, by exploiting polarisation from reflection under natural illumination (part III). The latter is, in general, partially linearly polarised. We start by giving a quick overview of the necessary mathematical background in polarisation (section 5.1) which we then apply to derive the equations for polarisation imaging under partially polarised incident illumination (section 5.2) and identify the conditions that allow us to link those expressions back to the well-studied case of polarisation imaging assuming unpolarised incident illumination. We provide practical guidelines for on-site acquisition (chapter 6) based on the theoretical analysis outlined in chapter 5 and demonstrate high quality results with an entry level DSLR camera (Sections 6.1 to 6.3) as well as with a mobile phone (section 6.4). Please note

that all rendered material in this dissertation is best appreciated on a screen.

Finally, chapter 7 discusses the overall contribution of the body of work presented in this thesis which are summarised as follows:

- We first present two novel mobile reflectometry approaches for acquiring detailed spatially-varying isotropic surface reflectance and mesostructure of planar material samples using a commodity mobile device, in general indoors environments (part II).
- We then present a novel approach for on-site acquisition of surface reflectance for planar, spatially-varying, isotropic samples in uncontrolled outdoor environments, which exploits the naturally occurring linear polarisation of incident and reflected illumination (part III).

We further put our work in perspective with recent research in the field of appearance acquisition and modelling, and discuss amenities for possible future work directions.

The work presented in this dissertation gave rise to the following two journal publications:

1. J  r  my RIVIERE, Pieter PEERS and Abhijeet GHOSH.
Mobile Surface Reflectometry.
Computer Graphics Forum, 35(1):191-202, 2016 [1].
Presented at Eurographics, May 2016.
2. J  r  my RIVIERE, Ilya RESHETOUSKI, Luka FILIPI and Abhijeet GHOSH.
Polarization imaging reflectometry in the wild.
ACM Transaction On Graphics (TOG), 36(6):206, 2017 [2].
Proceedings of ACM SIGGRAPH Asia 2017.

We also provide supplemental material in the form of videos and additional renderings under environmental illumination:

- Appendix A presents additional renderings under environmental illumination from the reflectance maps obtained with the mobile surface reflectometry

approaches presented in part II. Please also see the accompanying video showing our measurement protocols in action as well as animated renderings, available at: <https://youtu.be/vV291X0zfyU>.

- Appendix B presents additional renderings under environmental illumination from the reflectance maps obtained with the polarimetric reflectance method under natural outdoor illumination presented in part III. Please also see the accompanying video for animated results under environmental illumination at: <https://youtu.be/3xP6Z8s8AMQ>.

Chapter 2

Background And Related Work

In this chapter, we review some related work on reflectance acquisition and representation. We first introduce the concept of the Bidirectional Reflectance Distribution Function (BRDF) and its many properties and review some prior work on reflectance representation for both specular (Sections 2.1.1 to 2.1.3) and diffuse reflection (section 2.1.4). We then present prior art on BRDF acquisition, from the early ages in the field of reflectance acquisition, utilising complex setups restricted to the laboratory (section 2.2.1), to more accessible setups based on off-the-shelf hardware (section 2.2.2), similar in spirit to our mobile reflectometry approaches (part II). A recent in-depth survey on the topic of BRDF representation and acquisition can be found in [15] and [16].

We further review previous work on passive reflectometry under uncontrolled and/or unknown illumination (section 2.2.3), and finish our related work section by covering prior art on polarisation imaging in vision and graphics, with a particular emphasis on reflectance (section 2.3.1) and shape (section 2.3.2) estimation from polarisation. In our work on outdoors reflectometry (part III), we exploit polarisation cues from multiple views under outdoors natural illumination to resolve high-resolution reflectance maps of multiple planar surfaces that could not be acquired other than on-site.

2.1 The Bidirectional Reflectance Distribution Function (BRDF)

The main goal of reflectometry is to quantify the complex interaction of light with a material in order to obtain a compact representation for use in rendering applications. Nicodemus [7] was the first to formally define such interaction in terms of the Bidirectional Reflectance Distribution Function (BRDF), a 4D real-valued function defined over the hemisphere surrounding a differential surface area (dA). Figure 2.1 defines the geometry of the BRDF; given a pair of incident and outgoing directions (respectively $\vec{\omega}_i = (\theta_i, \phi_i)$ and $\vec{\omega}_o = (\theta_o, \phi_o)$), the BRDF returns the ratio of reflected radiance along $\vec{\omega}_o$ to the irradiance incident on the surface at $\vec{\omega}_i$.




Figure 2.1: **BRDF geometry:** The BRDF is a 4D real-valued function defined over the upper hemisphere (Ω^+) oriented by the normal to the surface (\vec{n}).

Mathematically, the BRDF for an homogeneous material is defined as a 4D function of incoming light direction ($\vec{\omega}_i = (\theta_i, \phi_i)$) and outgoing direction ($\vec{\omega}_o = (\theta_o, \phi_o)$) over the upper hemisphere ($\Omega^+ = [0, \frac{\pi}{2}] \times [0, 2\pi]$) oriented by the surface

normal \vec{n} :

$$f_r(\vec{\omega}_i, \vec{\omega}_o) = \frac{dL_o(\vec{\omega}_o)}{dE_i(\vec{\omega}_i)} \quad (2.1)$$

where

$$dE_i(\vec{\omega}_i) = L_i(\vec{n} \cdot \vec{\omega}_i) d\vec{\omega}_i$$

In eq. (2.1), $dE_i(\vec{\omega}_i)$ is the irradiance (i.e the incident flux of radiance per unit area along $\vec{\omega}_i$), $dL_o(\vec{\omega}_o)$ the reflected radiance (i.e the flux of outgoing radiance per unit area along $\vec{\omega}_o$) and \vec{n} is the surface normal, which defines the local orientation of the differential patch dA . For non-homogeneous materials, it is necessary to add two more degree of freedom to the BRDF, namely the position (\mathbf{x}) of the surface point. The surface is then said to exhibit a Spatially-Varying BRDF (SVBRDF). BRDFs are reciprocal and energy conserving:

$$\forall (\vec{\omega}_i, \vec{\omega}_o) \in (\Omega^+)^2, f_r(\vec{\omega}_i, \vec{\omega}_o) = f_r(\vec{\omega}_o, \vec{\omega}_i) \quad (\text{Helmholtz reciprocity}) \quad (2.2)$$

$$\forall \vec{\omega}_i \in \Omega^+, \int_{\Omega^+} f_r(\vec{\omega}_i, \vec{\omega}_o) (\vec{n} \cdot \vec{\omega}_o) d\vec{\omega}_o \leq 1 \quad (\text{Energy conservation})$$



Figure 2.2: **Diffuse vs Specular reflection:** Notice how the diffuse material looks identical as the camera moves while the specular material changes rapidly as the camera is tilted to the right.

Reflection can be separated in two categories. Diffuse reflection which exhibits slow to no view-dependent variations depending on the extent to which the incident light is being scattered before reaching the observer. The second category, specular reflection, is strongly view-dependent and corresponds to the shine at the surface of a material (see fig. 2.2). Given representations of specular ($f_s(\vec{\omega}_i, \vec{\omega}_o)$, Sections 2.1.1 to 2.1.3) and diffuse ($f_d(\vec{\omega}_i, \vec{\omega}_o)$, section 2.1.4) reflection models, a material's BRDF may be expressed according to the dichromatic reflectance model

[6] as:


$$f_r(\vec{\omega}_i, \vec{\omega}_o) = f_d(\vec{\omega}_i, \vec{\omega}_o) + f_s(\vec{\omega}_i, \vec{\omega}_o) \quad (2.3)$$

When a material behaves isotropically (i.e its BRDF is invariant by rotation around the surface normal), the BRDF can be reduced to a 3D function ($f_r(\theta_i; \theta_o; \phi_i - \phi_o)$), allowing for a dense representation as a tabulated 3D texture. However, in general, the BRDF can depend on many more parameters such as surface point position (\mathbf{x}), wavelength (λ), etc. Researchers have thus developed analytic models for compact BRDF representation, requiring only a few parameters which can be manually adjusted or fitted to measured data. In particular, in the remainder of this thesis, we will be interested in BRDF models for isotropic, spatially-varying reflection.

2.1.1 Fresnel reflectance

Fresnel reflectance accounts for the increase in specular reflection as the incident direction $\vec{\omega}_i$ approaches grazing angle (i.e. $\theta_i \rightarrow \frac{\pi}{2}$). As an example, the glare observed from a boat at the surface of a lake is due to Fresnel effects from reflection of the sun at grazing angle on the water. Augustin-Jean Fresnel, a French engineer and physicist, first derived the equations to quantify the behaviour of light at the interface between two media with differing refractive indices (η). For the purpose of this thesis, we will only consider Fresnel effects at a dielectric-dielectric interface (see fig. 2.3). Details of the dielectric-conductor equations can be found in [17].

Fresnel reflectance depends on the polarisation of the incident light and changes its polarisation state upon reflection. Mathematically, Fresnel effects at a



(a) Fresnel geometry: $\vec{\omega}_i$ is the incident direction and $\vec{\omega}_t$ the direction of transmission
 (b) Fresnel equations plotted for an air-glass interface ($\eta_i = 1$, $\eta_t = 1.5$)

Figure 2.3: **Fresnel visualisation:** Geometry (a) and equations at a dielectric-dielectric interface (b).

dielectric-dielectric interface are defined as:

$$\begin{aligned}
 R_{\perp}(\eta_i, \eta_t; \theta_i) &= \frac{\eta_i \cos \theta_i - \eta_t \sqrt{1 - \left(\frac{\eta_i}{\eta_t} \sin \theta_i\right)^2}}{\eta_i \cos \theta_i + \eta_t \sqrt{1 - \left(\frac{\eta_i}{\eta_t} \sin \theta_i\right)^2}} && \text{Reflectance of s-polarized light} \\
 R_{\parallel}(\eta_i, \eta_t; \theta_i) &= \frac{\eta_i \sqrt{1 - \left(\frac{\eta_i}{\eta_t} \sin \theta_i\right)^2} - \eta_t \cos \theta_i}{\eta_i \sqrt{1 - \left(\frac{\eta_i}{\eta_t} \sin \theta_i\right)^2} + \eta_t \cos \theta_i} && \text{Reflectance of p-polarized light}
 \end{aligned}
 \tag{2.4}$$

The unpolarised Fresnel reflection coefficient is obtained as the average of the two polarised coefficients:

$$F(\theta_i) = \frac{R_{\perp}(\theta_i) + R_{\parallel}(\theta_i)}{2}
 \tag{2.5}$$

In general, Fresnel equations are not used in the form of eq. (2.4), as they are rather complex and expensive to evaluate. Instead, researchers have proposed

cheaper alternatives in particular for real-time rendering applications. Cook and Torrance [18] proposed a simpler and exact formulation under the assumption of unpolarised incident illumination:

$$F(\theta_i) = \frac{1(g - \cos \theta_i)^2}{2(g + \cos \theta_i)^2} \left(1 + \frac{(\cos \theta_i(g + \cos \theta_i) - 1)^2}{(\cos \theta_i(g - \cos \theta_i) + 1)^2} \right) \quad (2.6)$$

where $g = \sqrt{\eta_t^2 - 1 + \cos^2 \theta_i}$.

Schlick [19] later proposed a polynomial approximation to eq. (2.5) for modelling unpolarised Fresnel effects which is often used in game development for its simplicity and cheap run-time evaluation:

$$F(\theta_i) = F(0^\circ) + (1 - F(0^\circ))(1 - \cos \theta_i)^5 \quad (2.7)$$

where $F(0^\circ)$ is the reflectance at normal incidence, which can be computed from the indices of refraction as:

$$F(0^\circ) = \left(\frac{\eta_i - \eta_t}{\eta_i + \eta_t} \right)^2 \quad (2.8)$$

2.1.2 Specular reflection models

Over the years, researchers have proposed multiple models to represent specular reflection, which can be classified in two categories:

1. Empirical BRDF models that are not physically accurate but provide a basis for reflectance estimation.
2. Physically-based BRDF models that derive an accurate representation of a given class of reflection based on the underlying physical properties of the material.

Empirical BRDF models

One of the very first empirical models was the Phong BRDF model [20], where the specular reflection is represented as a cosine lobe:

$$f_s(\vec{\omega}_i, \vec{\omega}_o) = \rho_s(\vec{\omega}_r \cdot \vec{\omega}_o)^s \quad (2.9)$$

where $\vec{\omega}_r = 2(\vec{\omega}_i \cdot \vec{n})\vec{n} - \vec{\omega}_i$ is the reflection vector obtained by reflecting the incident direction $\vec{\omega}_i$ about the surface normal \vec{n} , s is the shininess parameter which controls the size of the lobe and hence how shiny or dull the material appears and ρ_s is the specular albedo. This model is neither reciprocal nor energy conserving but can be slightly modified to yield energy conservation [21, 22]:

$$f_s(\vec{\omega}_i, \vec{\omega}_o) = \frac{s+2}{2\pi} \rho_s(\vec{\omega}_r \cdot \vec{\omega}_o)^s \quad (2.10)$$

Based on Phong's work, Blinn [23] proposed a more physically accurate BRDF which is the default in the fixed-function pipelines of OpenGL[®] and DirectX[®]. He noted that the highest reflection occurs when the halfway vector $\vec{\omega}_h = \frac{\vec{\omega}_i + \vec{\omega}_o}{|\vec{\omega}_i + \vec{\omega}_o|}$ is aligned with the surface normal and defined his model as:

$$f_s(\vec{\omega}_i, \vec{\omega}_o) = \rho_s(\vec{\omega}_h \cdot \vec{n})^s \quad (2.11)$$

Again, Blinn's original formulation is not energy conserving, but can be modified to ensure energy conservation [24]:

$$f_s(\vec{\omega}_i, \vec{\omega}_o) = \frac{(s+2)(s+4)}{8\pi(s+2^{-s/2})} \rho_s(\vec{\omega}_h \cdot \vec{n}) \quad (2.12)$$

In 1992, Ward [9] proposed to model anisotropic specular reflection with an elliptical Gaussian distribution of halfway vectors that is both reciprocal and energy conserving. This model is defined in a local tangent basis $[\vec{t}, \vec{b}, \vec{n}]$ where \vec{t}

and \vec{b} are the surface's tangent and bi-tangent respectively, such that $\vec{b} = \vec{t} \times \vec{n}$:

$$f_s(\vec{\omega}_i, \vec{\omega}_o) = \rho_s \frac{1}{\sqrt{\cos \theta_i \cos \theta_o}} \frac{\exp \left[-\tan^2 \theta_h \left(\frac{\cos^2 \phi_h}{\sigma_b^2} + \frac{\sin^2 \phi_h}{\sigma_t^2} \right) \right]}{4\pi \sigma_b \sigma_t} \quad (2.13)$$

where θ_h and ϕ_h are the elevation and azimuth angles of $\vec{\omega}_h$ with respect to the local tangent space, σ_b and σ_t are the surface's roughness along the tangent and bi-tangent respectively. Note that when $\sigma_b = \sigma_t = \sigma$, eq. (2.13) reduces to an isotropic BRDF:

$$f_s(\vec{\omega}_i, \vec{\omega}_o) = \rho_s \frac{1}{\sqrt{\cos \theta_i \cos \theta_o}} \frac{\exp \left[\frac{-\tan^2 \theta_h}{\sigma^2} \right]}{4\pi \sigma^2} \quad (2.14)$$

More recently, Ashikmin and Shirley [25] proposed a reciprocal and energy conserving anisotropic model based on Phong's lobe (eq. (2.9)), further modelling Fresnel reflectance effects (section 2.1.1):

$$f_s(\vec{\omega}_i, \vec{\omega}_o) = \frac{\sqrt{(s_b + 1)(s_t + 1)} (\vec{n} \cdot \vec{\omega}_h)^{(s_b \cos^2 \phi_h + s_b \sin^2 \phi_h)} F(\theta)}{8\pi (\vec{\omega}_h \cdot \vec{\omega}_i) \max\{(\vec{n} \cdot \vec{\omega}_i), (\vec{n} \cdot \vec{\omega}_o)\}} \quad (2.15)$$

where s_t and s_b are the shininess exponents along the tangent and bi-tangent respectively, $F(\theta)$ accounts for Fresnel reflectance and θ is the angle between $\vec{\omega}_h$ and $\vec{\omega}_i$.

Physically-based BRDF models (Microfacet theory)

Physically-based BRDF models were first proposed in the optics literature and quickly adopted in graphics as the de-facto models for photo-realistic rendering. Torrance and Sparrow [26] were the first to propose a BRDF model based on the microfacet theory for applications in computer graphics. They derived their model under the assumption that the surface is composed of a collection of perfectly specular micro surfaces, statistically distributed around the normal to the surface

(\vec{n}) and their model was later improved by Cook and Torrance [18]:

$$f_s(\vec{\omega}_i, \vec{\omega}_o) = \frac{F(\theta)G(\vec{\omega}_i, \vec{\omega}_o, \vec{n})D(\vec{\omega}_i, \vec{\omega}_o, \vec{n}, \sigma)}{\pi|\vec{n} \cdot \vec{\omega}_i||\vec{n} \cdot \vec{\omega}_o|} \quad (2.16)$$

where $F(\theta)$ models Fresnel reflectance effects (section 2.1.1) with θ the angle between $\vec{\omega}_i$ and $\vec{\omega}_h$, $G(\vec{\omega}_i, \vec{\omega}_o, \vec{n})$ is the shadowing-masking term which accounts for the fact that any microfacet may either be shadowed or masked by any neighbouring facet and $D(\vec{\omega}_i, \vec{\omega}_o, \vec{n}, \sigma)$ is the distribution term which models the statistical distribution of microfacet normals ($\vec{\omega}_h$) about the surface normal (section 2.1.3). Recent papers [27, 28] disagree with the π term in the denominator and instead propose to rewrite eq. (2.16) as:

$$f_s(\vec{\omega}_i, \vec{\omega}_o) = \frac{F(\theta)G(\vec{\omega}_i, \vec{\omega}_o, \vec{n})D(\vec{\omega}_i, \vec{\omega}_o, \vec{n}, \sigma)}{4|\vec{n} \cdot \vec{\omega}_i||\vec{n} \cdot \vec{\omega}_o|} \quad (2.17)$$

2.1.3 Microfacet distribution models

The microfacet distribution term ($D(\vec{\omega}_i, \vec{\omega}_o, \vec{n}, \sigma)$) is a bell-shaped curve quantifying the statistical distribution of microfacet normals ($\vec{\omega}_h$) about the normal to the surface (\vec{n}). Its width is controlled by a roughness parameter σ which is the Root Mean Squared (RMS) slope of the microfacets. The bidirectional shadowing-masking term $G(\vec{\omega}_i, \vec{\omega}_o, \vec{n})$ depends on both the distribution $D(\vec{\omega}_i, \vec{\omega}_o, \vec{n}, \sigma)$ and the underlying structure of the micro surfaces and accounts for self-occlusion as well as occlusion by neighbouring microfacets.

Torrance and Sparrow and later Cook and Torrance proposed to model the distribution of microfacet normals ($\vec{\omega}_h$) about the surface normal by a Beckmann distribution term [29], which is a Gaussian lobe defined as:

$$\begin{aligned} D(\vec{\omega}_i, \vec{\omega}_o, \vec{n}, \sigma) &= \frac{1}{\pi\sigma^2 \cos^4 \theta_h} e^{-\frac{\tan^2 \theta_h}{\sigma^2}} \\ &= \frac{1}{\pi\sigma^2 \cos^4 \theta_h} e^{-\frac{(\vec{\omega}_h \cdot \vec{n})^2 - 1}{(\vec{\omega}_h \cdot \vec{n})^2 \sigma^2}} \end{aligned} \quad (2.18)$$

In their original paper, Cook and Torrance proposed to model the shadowing-masking term $G(\vec{\omega}_i, \vec{\omega}_o, \vec{n})$ assuming a V-groove shape for the microfacets (see [23, 30, 26] for more details):

$$G(\vec{\omega}_i, \vec{\omega}_o, \vec{n}) = \min \left\{ 1, \frac{2(\vec{\omega}_h \cdot \vec{n})(\vec{\omega}_o \cdot \vec{n})}{(\vec{\omega}_h \cdot \vec{\omega}_o)}, \frac{2(\vec{\omega}_h \cdot \vec{n})(\vec{\omega}_i \cdot \vec{n})}{(\vec{\omega}_h \cdot \vec{\omega}_o)} \right\} \quad (2.19)$$

However, Walter et al. [27] recently recommended to use the Smith shadowing-masking term instead [31], in particular for data fitting applications, to circumvent the non differentiability of eq. (2.19). The Smith shadowing-masking term is defined as a separable product of two mono-directional terms G_1 :

$$G(\vec{\omega}_i, \vec{\omega}_o, \vec{n}) \approx G_1(\vec{\omega}_i, \vec{\omega}_h)G_1(\vec{\omega}_o, \vec{\omega}_h) \quad (2.20)$$

where G_1 is derived directly from the distribution term [31, 32, 33, 27]. For the Beckmann distribution, G_1 has an analytic solution to which Schlick [19] proposed a polynomial approximation:

$$G_1(\vec{\omega}, \vec{\omega}_h) = \frac{2}{1 + \operatorname{erf}(a) + \frac{1}{a\sqrt{\pi}e^{-a^2}}} \approx \begin{cases} \frac{3.535a + 2.181a^2}{1 + 2.276a + 2.577a^2} & \text{if } a < 1.6 \\ 1 & \text{otherwise} \end{cases} \quad (2.21)$$

where $a = \frac{1}{\sigma \tan \theta_h}$

In the same paper, Walter et al. also proposed a new distribution term to model refraction through rough surfaces. Their distribution, which they named GGX, is based on the Trowbridge-Reitz distribution function [34] and is defined

mathematically (with its associated G_1 term) as:

$$D(\vec{\omega}_i, \vec{\omega}_o, \vec{n}, \sigma) = \frac{\sigma^2}{\pi \cos^4 \theta_h (\sigma^2 + \tan^2 \theta_h)^2} \quad (2.22)$$

$$G_1(\omega, \vec{\omega}_h) = \frac{2}{1 + \sqrt{1 + \sigma^2 \tan^2 \theta}}$$

The GGX distribution term has rapidly become the preferred model for physically-based rendering of specular highlights from rough surfaces and is the model we employ for data fitting and rendering in the remainder of this thesis, together with the expression of the microfacet BRDF presented in eq. (2.17).

distribution-based BRDF (dBRDF)

In 2007, Ashikhmin and Premože proposed a generalization of the Ashikhmin-Shirley anisotropic Phong BRDF model (eq. (2.15)) to enable the use of any microfacet distribution while having a simpler mathematical form compared to Ward's anisotropic model (eq. (2.13)):

$$f_s(\vec{\omega}_i, \vec{\omega}_o) = \rho_s \frac{D(\vec{\omega}_i, \vec{\omega}_o, \vec{n}, \sigma) F(\theta)}{(\vec{\omega}_i \cdot \vec{n}) + (\vec{\omega}_o \cdot \vec{n}) - (\vec{\omega}_i \cdot \vec{n})(\vec{\omega}_o \cdot \vec{n})} \quad (2.23)$$

where they model Fresnel effects with Schlick's approximation eq. (2.7) with θ the angle between $\vec{\omega}_h$ and $\vec{\omega}_i$.

The main advantage of this model is that the distribution may be extracted directly from measured data without requiring numerical fitting simply by taking measurements in the back-scattering direction (i.e when $\vec{\omega}_i = \vec{\omega}_o = \vec{\omega}_{bs}$). Equation (2.23) can then be simplified as:

$$f_r(\vec{\omega}_{bs}) = \frac{\rho_s F(0^\circ) D(\vec{\omega}_{bs}, \vec{n}, \sigma)}{2(\vec{\omega}_{bs} \cdot \vec{n}) - (\vec{\omega}_{bs} \cdot \vec{n})^2} \quad (2.24)$$

This expression informs us that measurements made in the back-scattering direction are directly proportional to the distribution of microfacets which we will exploit in chapter 3 to fit SVBRDFs from data obtained with a mobile device in

the back-scattering direction.

2.1.4 Diffuse reflection models

Diffuse reflection accounts for the portion of light that is scattered multiple times inside the material, before being reflected back out towards the observer. Visually, it corresponds to the matte texture of the material. Depending on the number of scattering events occurring within the material, scattered light may exhibit some low-frequency view dependency which can be explained by subsurface scattering models. These however go beyond the scope of this thesis, for which we assume diffuse reflection to stem from pure Lambertian reflection.

The simplest BRDF for diffuse reflection is the Lambertian model, which models a perfect diffusely reflecting smooth surface as a constant BRDF:

$$f_d(\vec{\omega}_i, \vec{\omega}_o) = \frac{\rho_d}{\pi} \quad (2.25)$$

where ρ_d is the surface's diffuse albedo.

In order to account for diffuse reflection from rough surfaces, Oren and Nayar [35] proposed a generalisation of the Lambertian model based on the microfacet theory (section 2.1.2). They model a rough diffuse surface as a collection of perfect Lambertian microfacets:

$$f_d(\vec{\omega}_i, \vec{\omega}_o) = \frac{\rho_d}{\pi} (A + (B \cdot \max\{0, \cos(\phi_i - \phi_o)\} \cdot \sin \alpha \cdot \tan \beta)) \quad (2.26)$$

where $A = 1 - 0.5 \frac{\sigma^2}{\sigma^2 + 0.33}$, $B = 0.45 \frac{\sigma^2}{\sigma^2 + 0.09}$, $\alpha = \max\{\theta_i, \theta_o\}$, $\beta = \min\{\theta_i, \theta_o\}$ and σ is the surface's roughness.

2.2 Reflectometry

Physically-based BRDF models alone are generally not sufficient to produce compelling, photo-realistic renderings, as they do not account for subtle details such as scratches, wearing and rusting, which are important to convey an impression of realism. Arguably the best way to reproduce the rich variation in reflectance of real-world materials is to measure said materials, and extract their intrinsic physical properties, in order to drive a physically-based rendering engine. Researchers have therefore looked at designing acquisition setups and algorithms for measuring the BRDF of materials. In this section, we will review some related work on reflectance acquisition, covering the very first designs for dense reflectometry (section 2.2.1) which require expensive hardware only suited to a laboratory environment, followed by setups based on commodity hardware and free-form acquisition (section 2.2.2) and approaches for passive reflectometry (section 2.2.3), which somewhat relax the need to carry measurements in a controlled environment.

2.2.1 Laboratory setup

Dense measurement

One of the very first designs for reflectometry, the gonireflectometer (fig. 2.4), consists of a pair of reflectance detector and point light source mounted on separate mechanical gantries, allowing to cover the entire hemisphere of directions (Ω^+) around the sample under study. This design was used by many researchers to build accurate databases of densely measured BRDFs for various material samples [8, 36].

While extremely accurate, measurements with a gonireflectometer suffer from lengthy run-times due largely to the use of a reflectance detector. Furthermore, the density at which the measurements can be made greatly depends on the mechanical design of the reflectometer which generally hinders measurement at grazing angle and in the back-scattering direction, thus requiring extrapolation




Figure 2.4: **Gonioreflectometer:** Schematic of a typical gonioreflectometer as proposed by Murray-Colemann and Smith [3].

to account for any missing data.

Image-based measurement

In order to reduce measurement times, many researchers have proposed revised designs to the original gonioreflectometer. In particular, thanks to the rapid development in DSLR camera technology, image-based measurement apparatus have rapidly become the standard in reflectometry. In 1992, Ward [9] proposed to employ a fish-eye lens mounted to a Charge-Coupled Device (CCD) camera coupled with a half-silvered dome, in order to measure a 2D slice of the BRDF from all possible viewing directions at once. Another common way to measure every possible viewing directions at once is by photographing a convex object (generally a sphere) made of the material under study [10]. Following this idea, Matusik et al. [37] proposed an image-based setup for dense BRDF measurement of homogeneous spheres, and used their setup to build the MERL database, composed of 100 different materials.

2.2.2 Accessible reflectometry

Although extremely accurate, the previous setups are expensive and require a wealth of knowledge in reflectometry to be handled correctly thus limiting their use to experts in the field. Recently, with the development of mobile technology and with digital photography equipment becoming more and more accessible, a

new trend in research has emerged, to address the accessibility of reflectometry to non-expert users. This new trend can be roughly classified in two categories: acquisition using off-the-shelf hardware and free-form acquisition. The main aim of both approaches is to simplify the capture process so as to make reflectometry accessible to digital artists who would otherwise have to go through the time consuming process of hand crafting physically plausible reflectance maps.

Commodity hardware

No matter how easy or complex a measurement setup is for reflectometry, it requires a reflectance detector or camera and light source at the bare minimum. Most often, researchers have proposed to employ a camera in conjunction with a linear light source, where the linear light source helps reduce the amount of data required, especially for SVBRDF acquisition. Gardner et al. [38] built a mechanical gantry out of LEGOTM bricks to which they attached a linear light source. By translating the gantry above a planar surface and recording a video of the sample at a fixed vantage point, they were able to capture spatially-varying reflectance maps of planar surfaces. This design was later modified by Chen et al. [39] to measure anisotropic BRDFs.

Another common setup for acquisition with commodity hardware is to pair a DSLR camera with an LCD monitor, to provide an extended source of illumination, further reducing time spent in measurements. Francken et al. [40] recently proposed employing Gray code patterns projected from an LCD panel to acquire normals and Phong exponent maps of glossy surfaces. Also targeted at glossy surfaces, the approach of Wang et al. [41] uses step-edge illumination to measure specular reflectance and bump maps. Ghosh et al. [5] proposed to project spherical gradient illumination patterns from an LCD monitor to recover model-independent reflectance parameters, exploiting the inherent polarisation of the LCD panel for diffuse-specular separation. More recently, Aittala et al. [42] have proposed to capture a sample's response to 2D Fourier patterns in order to resolve a complete set of per-pixel reflectance maps, namely: diffuse and specular albedos, normals

and roughness. Our work on mobile reflectometry using an LCD-based setup (chapter 4) is closest in spirit to the latter two papers. The main differences stem from the fact that our method is self-contained and only requires that a small representative patch of the sample be measured.

Free-form acquisition

Free-form acquisition designs have become popular over the past 15 years due largely to their practicality and ease of use: one or both of the image device and light source are operated by hand so as to sample a material's reflectance properties. Masselus et al. [13] first proposed the free-form light-stage, where they move a hand-held calibrated light source around the object to record its reflectance field from a fixed view point for use in image-based relighting. Drawing inspiration from both Masselus et al. [13] and Gardner et al. [38], Ren et al. [12] proposed to use a hand-waved linear light source coupled with a mobile phone camera and custom-made BRDF chart, where the latter provides a basis for BRDF fitting.

With the current advances in mobile technology, researchers have started to investigate more compact and portable designs for reflectance acquisition. RGB-D cameras such as Microsoft®'s KinectTM motion sensing device, which was first created as a controller for the Xbox 360 gaming system, rapidly became popular in the research community as a tool for material scanning. Wu and Zhou [43] proposed an integrated system for hand-held acquisition of shape and reflectance based on a Kinect sensor. Furthermore, a wealth of commercial applications and open source projects have been developed around Microsoft's Kinect technology ranging from facial reconstruction to entire rooms scanning ([44, 45]).

Close to our work (chapter 3), is that of Aittala et al. [14] who recently proposed a two-shot method for the acquisition of stochastic materials using a mobile phone. They employ a pair of flash-no flash observations of the sample in the back-scattering direction, in general indoor environments, coupled with statistical analysis to extract reflectance maps of planar stochastic surfaces. They later extended their method to a single flash image for stationary materials using a

deep learning approach for texture synthesis [46]. Recently, Li et al. [47] proposed a method to estimate spatially-varying reflectance properties from a single photograph using self-augmented Convolutional Neural Networks (CNNs). As with many CNN-based approaches, the main drawback of these methods lies in the fact that they require a large amount of labelled BRDF data and a different network per class of material. Unlike these methods, our free-form acquisition method (chapter 3) is more general, as it does not require the sample to be stochastic to resolve detailed, spatially-varying reflectance maps, and requires relatively little data (around 250 frames in most of our examples). Furthermore, unlike Attaila et al. [14, 46], our method produces an exact depiction of the acquired material, rather than a statistically plausible representation of the material’s appearance.

2.2.3 Uncontrolled environments

Reflectometry under uncontrolled and/or unknown lighting is an extremely challenging problem that has attracted the attention of researchers both in vision and graphics. In 2008, Glencross et al. [48] proposed a flash-no flash method for depth hallucination to recover depth and spatially-varying diffuse albedo of planar surfaces acquired outdoors. While producing perceptually plausible results under complex outdoors illumination, their method is limited to diffuse samples only under diffuse illumination. The same year, Romeiro et al. [49] proposed an image-based approach to reflectometry under passive, uncontrolled illumination, assuming curved objects of homogeneous properties, which they later extended to unknown illumination, leveraging the statistics of natural illumination [50]. Lombardi and Nishino [51] later proposed a method to estimate reflectance and lighting for objects with known shape and homogeneous material properties, from a single photograph. They employ an expectation maximisation approach with appropriate priors on both the BRDF (directional statistics) and illumination (natural image statistics). A similar approach has also been employed by Oxholm and Nishino, for the joint estimation of shape and homogeneous BRDF under known illumination [52]. With the exception of the method of Glencross et al., the main limitation of

the previous methods is that they require the material to be homogeneous, and often times assume a spherical object.

Many researchers have also investigated methods to resolve SVBRDFs under uncontrolled illumination. Similar to Glencross et al., Hauagge et al. [53] also assume a Lambertian material and a model of unoccluded sun-sky illumination, to recover per-pixel diffuse albedo from collections of images taken from the Internet. More recently, Dong et al. [54] further recovered spatially varying isotropic reflectance properties from a video of a rotating object (with known shape) under unknown illumination, which they coined “Appearance from Motion”. Their approach alternatively estimates reflectance and lighting in an iterative process. Xia et al. [55] further extended this approach for simultaneously estimating reflectance, object shape and illumination .

2.3 Polarisation in vision and graphics

Polarisation plays an important role in vision and graphics as it provides useful cues for reflectance estimation, material classification and shape estimation. Most prior art on reflectance and/or shape estimation from polarisation has studied polarisation from reflection due to purely linearly polarised incident illumination [56, 57, 58] or unpolarised incident illumination [59, 60]. Two notable exceptions are Koshikawa et al. [61] and Ghosh et al. [62] who proposed to exploit circular polarisation for shape and reflectance estimation respectively. In this section, we will review some previous work on reflectometry based on polarisation without going into the mathematical details of the underlying physics. Instead, these will be covered in depth in chapter 5 where we will lay the foundations for polarisation imaging in terms of Mueller calculus (section 5.1). We will then apply this theory to the previously unstudied case of polarisation imaging under partially linearly polarised incident illumination (section 5.2) which we show to be a generalisation of the well-studied approach of polarisation imaging under unpolarised incident illumination. This will constitute the theoretical foundations on which we build

our method for passive reflectometry in the wild (chapter 6).

2.3.1 Reflectance separation/estimation

Most BRDF models follow the dichromatic model [6], i.e they represent reflection as a linear combination of a diffuse and specular component (eq. (2.3)), which can be extremely difficult to resolve jointly. One of the most important pre-requisite to appearance modelling is therefore to accurately separate surface reflectance into its diffuse and specular components. Many researchers have looked at polarisation imaging for this purpose because of the intrinsic property of reflection on polarisation: diffuse reflection, which is the result of multiple scattering inside the material, tends to depolarise light, while specular reflection (due to a single scattering event at the surface) preserves the polarisation properties of polarised incident illumination. Under unpolarised incident illumination, a reflected ray of light becomes partially linearly polarised depending on the angle of incidence, due to Fresnel effects (see fig. 2.3b): as the angle of incidence varies away from normal incidence ($\theta_i = 0^\circ$), the reflectance of s-polarised (perpendicular polarisation) light increases while that of p-polarised (parallel polarisation) light decreases until it reaches 0 at a specific angle called Brewster angle (θ_B). At this specific angle the Degree of Polarisation (DOP) of the reflected beam reaches 100 %: the reflected beam becomes purely horizontally polarised with respect to the plane of incidence (see fig. 2.5, [blue curve](#) for a plot of the degree of specular polarisation at an air-glass interface).

Wolff and Boult were among the first researchers to consider a polarisation-aware reflectance model where they model both diffuse and specular polarisation resulting from Fresnel effects [63, 59], which they use for classifying materials into dielectrics and metals. Since then, many researchers have considered polarisation-aware reflectance models for appearance modelling. Compared to Wolff and Boult’s approach however, most approaches for polarisation-based appearance modelling have disregarded diffuse polarisation, as it generally accounts for less than 10% of the total reflected light [59, 64] away from occluding contours.




Figure 2.5: **Degree of Polarisation (DOP) \mathcal{P}** : Plots of the degree of specular polarisation (blue) and degree of diffuse polarisation (red) for unpolarised incident illumination reflected at an air-glass interface.

It is important to clearly distinguish diffuse reflection from diffuse polarisation: diffuse reflection corresponds to multiple subsurface scattering events which tend to depolarise light. On the other hand, diffuse polarisation accounts for polarisation resulting from unpolarised light (from diffuse reflection) being transmitted back out towards the observer, thus becoming partially polarised due to Fresnel effects (see fig. 2.5, **red curve** for a plot of the degree of diffuse polarisation at an glass-air interface).

Müller [60] proposed a method for diffuse-specular separation under the assumptions that the diffuse component completely depolarises light and the material under study is dielectric. The main practical limitation of his method is that it requires the index of refraction of the material to be known before-hand which in general is not the case. Ma et al. [57] proposed applying circular and linear polarisation with spherical gradient illumination (using a geodesic dome composed of 150 controllable LEDs) to obtain high quality diffuse and specular albedo and normal maps. Their method is however limited to a single view point. Ghosh et al. [65] later relaxed view-dependent pattern of Ma et al. by placing vertically polarised cameras at the equator of the dome and horizontal polarisers on the lines of latitude of the spherical illumination setup for diffuse, cross-polarised measurements. They further employed parallel polarised measurements (vertical on both cameras and lights) to observe mixed diffuse and specular reflection. In our LCD-based

approach (chapter 4), we also employ parallel and cross-polarised observations for diffuse-specular separation.

Close in spirit to our work on polarimetric reflectometry in-the-wild (chapter 6), is that of Miyazaki et al. [66] who employ polarisation imaging under the unpolarised world assumption (i.e the incident illumination is unpolarised), coupled with inverse rendering to jointly estimate the reflectance properties of homogeneous, convex objects and illumination conditions. Their method is however limited to rather simple scenes with few lights (3 in their examples) and requires the incident illumination to be unpolarised, which is generally not the case outdoors. Ghosh et al. [62] recently proposed a complete framework for recovering detailed spatially-varying reflectance properties of planar surfaces, by measuring the complete set of Stokes parameters resulting from reflection under circularly polarised incident illumination. Again, their method is not applicable outdoors as they require active illumination (spherical gradient illumination) for shape estimation and circular polarisation which does not occur naturally.

2.3.2 Shape estimation

Shape from polarisation has been thoroughly studied in the vision community, under the unpolarised world assumption: in these conditions, the polarisation state of the reflected ray of light provides cues to constrain the surface normal (\vec{n}) to the plane of incidence, which by definition contains the incident and outgoing directions ($\vec{\omega}_i$ and $\vec{\omega}_o$) and surface normal (\vec{n}). A common technique for polarisation imaging consists of observing a scene with a linear polariser at 3 or more orientations to measure the reflected radiance through the polariser which has the form of a phase-shifted sinusoid of phase ϕ (see fig. 2.6). Knowledge of the maximum and minimum intensities and phase of the sinusoid give constraints for shape and reflectance estimation. Chapter 5 will cover polarisation imaging in more depth. In this section we will only review some relevant previous material on Shape from Polarisation where much efforts have been focused on resolving the ambiguities inherent to polarisation imaging.

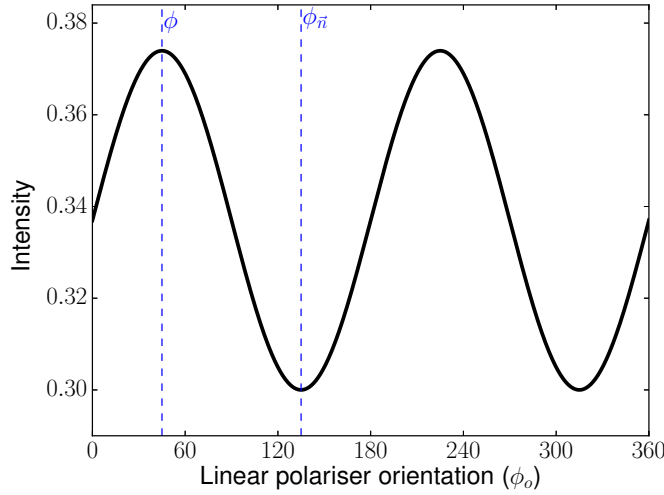


Figure 2.6: **Transmitted Radiance Sinusoid:** The observation of the radiance resulting from specular reflection through a rotating linear polariser has the form of a phase-shifted sinusoid, where the phase (ϕ) is related to the azimuth of the surface normal ($\phi_{\vec{n}}$) by $\pm \frac{\pi}{2}$.

Upon reflection on a specular surface, unpolarised incident illumination becomes partially polarised due to Fresnel effects, reaching maximum polarisation at Brewster angle (θ_B), where p-polarised light is completely transmitted (see fig. 2.3b). When observing reflected light through a linear polariser rotated at different orientations, the transmitted radiance has the form of a phase-shifted sinusoid (fig. 2.6) of phase ϕ . Wolff [67] was the first to consider the phase of the Transmitted Radiance Sinusoid (TRS) as a cue for shape estimation from multiple views: the main idea is that each view constrains the surface normal (\vec{n}) to lie within the plane of incidence, so in principle two views suffice to fully determine \vec{n} . The main advantage of this method is that, unlike other methods (e.g. Saito et al. [68]), it does not require that the index of refraction of the material be known before-hand, with the disadvantage of requiring per-pixel correspondence across views. This approach was later refined by Miyazaki et al. [69] to estimate the shape of transparent objects and Sadjadiz and Sadjadi [70] for shape and index of refraction estimation in the infrared domain. In our work (chapter 6) we also follow a multi-view Shape from Polarisation (SfP) approach, and propose practical guidelines to relax the unpolarised world assumption, based on a theoretical analysis of polarisation imaging in the previously unstudied case of partially linearly

polarised incident illumination (chapter 5).

Saito et al. [68] also rely on the TRS for estimating the shape of transparent objects. They fix the light and camera and rotate the object on a calibrated turn-table while densely sampling the TRS (36 samples per rotation step of the turn-table). From the maximum and minimum values of the TRS, they indirectly measure the reflected Degree of Polarisation (DOP) which is related to the normal's zenith angle ($\theta_{\vec{n}}$) by a one-to-one non linear mapping, while the azimuth angle ($\phi_{\vec{n}}$) is resolved from the phase (ϕ) of the TRS.

Guarnera et al. [71] follow a very similar approach to that of Saito et al. while requiring only 4 images in total for shape estimation. In their work, they measure the complete set of Stokes parameters (section 5.1 will cover Stokes formalism in details) under unpolarised as well as circularly polarised incident illumination, resolving the $\phi_{\vec{n}}$ -ambiguity in a similar way as Atkinson et al. [64]. The latter noted that the approach of shape from specular polarisation is ambiguous in general for $\theta_{\vec{n}}$, because the DOP reaches a maximum at Brewster angle (see fig. 2.5, [blue curve](#)). To circumvent this ambiguity, they instead proposed to measure the DOP due to diffuse reflection, which is monotonically increasing with respect to $\theta_{\vec{n}}$ (see fig. 2.5, [red curve](#)). In order to resolve the $\phi_{\vec{n}}$ -ambiguity, they proposed to populate the normals inwards from the object's silhouette, under the assumption of a concave geometrical shape.

Very recently, Kadambi et al. [72] proposed to resolve the $\phi_{\vec{n}}$ -ambiguity by augmenting coarse 3D scans with specular SfP under unpolarised incident illumination from a single view, where the coarse 3D geometry provides a mean to disambiguate the normals obtained from SfP. Smith et al. [73] have recently proposed direct inference of surface depth instead of normals by combining specular and diffuse polarisation cues with a linear depth constraint formulation. They demonstrate depth recovery under uncalibrated (unpolarised) point light sources as well as low order spherical harmonic illumination.

Prologue - Summary


With the recent advances in computer vision, graphics and the democratisation of affordable photography hardware, the field of reflectometry has come a long way towards digitizing real-world materials with a great degree of realism. Obtaining good digital clones however has so far been limited to controlled environments, thus limiting the scope of measurements to materials that can be physically brought into a laboratory. Furthermore, such laboratory setups generally require expensive equipment which require careful calibration by experts in the field.

Following the recent trend to make appearance modelling more affordable and user friendly by taking advantage of the availability of good quality, affordable hardware, we set out to further push the limits of reflectometry by investigating novel methods for on-site reflectometry based on off-the-shelf hardware in both semi-controlled environment (part II) and under passive illumination in outdoors conditions (part III).

PART II:

MOBILE SURFACE REFLECTOMETRY

[1]



(a) Free-form acquisition



(b) LCD-based acquisition



(c) Appearance transfer

Figure I: **Examples of surface reflectance recovered using Mobile Surface Reflectometry:** (a) A spatially varying rough specular material acquired using our hand-held free-form acquisition technique (chapter 3). (b) Highly specular surface reflectance recovered using mobile LCD-based reflectometry (chapter 4), with enhanced mesostructure from close-up observations under natural lighting (section 3.4). (c) Surface reflectance of a large spatially-varying material sample recovered using appearance transfer from surface reflectance obtained using the LCD-based approach for a small reference patch (section 4.3).

Chapter 3

Mobile Surface Reflectometry - Free-form Acquisition

Mobile devices nowadays are pervasive, with an estimated 11.6 billion connected mobile devices by 2020 according to Cisco [74], mostly used for mobile Internet access. However, with the advances in mobile technology, manufacturers are striving to pack their devices with the finest in mobile computing, photography, etc. In particular, every new mid-to-high-range phone or tablet out every few months is equipped with the latest technology in mobile digital photography. Of interest to us is the coupling of a back-facing high-resolution sensor, with an LED flash to allow quick and easy point and shoot photography.

Given the intrinsic nature of the back-facing camera/flash pair typically arranged to be near-coaxial, we designed a reflectance acquisition setup using a commodity mobile device, to measure the reflectance of spatially varying planar surfaces in the direction of back-scattering, suitable for fitting to any microfacet BRDF model [75]. In our mobile flash-based method, we propose an intuitive capture process, to allow quick and easy measurements: the user manually waves the mobile device pointed toward a static planar surface, in order to densely sample the upper hemisphere around the sample in the back-scattering direction (section 3.1). During the capture process, we show visual feedback to the user in real-time on the device's LCD screen.

A capture session only takes a few minutes and typically requires less than 250 frames captured at different orientations of the device. The data is then transferred to a Personal Computer (PC) for processing (section 3.3): we first calibrate the data both geometrically and radiometrically (section 3.2), and estimate the surface’s normals (section 3.3.1). We then proceed to separating the reflectance data into its diffuse (section 3.3.2) and specular components and estimate per-pixel roughness parameters for a GGX [27] BRDF model (section 3.3.3).

Section 3.3.4 presents some results of spatially-varying planar surfaces that we acquired with our proposed free-form acquisition method. While we are able to capture the general appearance of the material with this approach, some high frequency details are lost in the process due to the hand-held nature of the capture process. We thus propose a detail enhancement framework (section 3.4) to add back the missing high-frequency details, for more realism. Finally, in section 3.5, we discuss a few limitations of our method which we address in chapter 4.

3.1 Setup


In our experiments, we employed a Fujitsu Stylistic M532 10” Android tablet,



Figure 3.1: **Acquisition setup:** 10” Fujitsu Stylistic M532 Android tablet with an 8 MP back-facing camera with co-located LED flash.

with an 8 MP back facing camera coupled with a near co-axial LED flash (see

fig. 3.1) to capture reflectance data in the back-scattering direction at 1080p resolution. Besides the sample under study, we also require an X-Rite ColorChecker® calibration target to be visible in the scene for radiometric calibration of the acquired data. This is a direct consequence of some technical shortcomings of the technology at the time when we developed our method: our device runs Android 4.2 “Jelly Bean” which does not have support for HDR imaging¹. We thus instead rely on the presence of the calibration target in the captured data for radiometric calibration of the acquired data (section 3.2).



(a) **Data acquisition:** Notice the reflection of the half-opened shades on the device’s screen.

(b) The user hand waves the device pointed at the sample in order to densely measure its response in the back scattering direction

Figure 3.2: **Mobile Surface Reflectometry - free-from acquisition:** Flash-based surface reflectometry in a dimly lit office room.

The acquisition process proceeds as follows. The user points the mobile device’s back camera (and flash) at the planar sample from a distance of roughly 50 cm (fig. 3.2a) and densely samples the direction of back-scattering $\vec{\omega}_{bs}$ (fig. 3.2b) from multiple directions over the upper hemisphere. While most previous methods for reflectometry require a dark room, we only require that the flash’s illumination dominates ambient lighting. With our device, at a distance of roughly 50 cm, it corresponds to light levels of up to 30 to 40 LUX² (see reflections of the window from fig. 3.2a).

In all of our experiments, we start the capture process from normal incidence, in order to have a canonical view of reference to which we register every subsequent frame in the sequence. In theory, any frame from the capture sequence could be

¹The support for HDR imaging was introduced with Android 5 “Lollipop”

²This roughly corresponds to a dimly lit room.

used as a reference, but we pick the one at normal incidence as the extracted reflectance maps are then readily available in an axis-aligned texture space. We further require the presence of an X-Rite ColorChecker® for calibration of the hand-held measurements, to account for variations in distance and camera auto exposure.

In all of our experiments, we set the capture rate to 15 FPS³. At such frame-rate, a typical capture sequence corresponds to around 250 frames, recorded at multiple viewpoints. While in theory, we could use the calibration target for geometric calibration as well, in practice it is not always possible to keep the colour chart and sample fully visible for all frames. Instead we thus detect and track sparse 2D features [76] across the sequence of frames and warp every view to the canonical frame at normal incidence and propose two complementary methods for estimating the direction of back-scattering (section 3.2).

3.2 Calibration

The first step after acquiring the data is to calibrate it both geometrically and radiometrically. Geometric calibration is essential, as we require per-pixel alignment of the data as well as knowledge of the back-scattering direction in order to recover surface normals (section 3.3.1), as well as fit the data to a microfacet BRDF model (section 3.3.3).

3.2.1 Data registration

In order to obtain per-pixel correspondence throughout the captured sequence of frames, we track a sparse set of salient corners [76] by computing optical flow [77] between pairs of consecutive frames, thus obtaining a sparse set of features x_0 for the first frame in the sequence (canonical pose) and corresponding features x_i in each subsequent frame. Given these 2D correspondences in image plane, we

³This is the limit for frame grabbing at full High Definition (HD) resolution on our device.

compute the homography H_i such that:

$$x_0 = H_i x_i \quad (3.1)$$

and warp each frame of the sequence to align to the canonical frame. To overcome the problem of drift that occurs in flow-based tracking, we further adopt a reset mechanism that bootstraps the corner detection and tracking from a novel starting frame every 10 frames. Furthermore, our homography computation is made robust to any outliers present after the matching process, though the use of the RANdom SAmple Consensus (RANSAC) algorithm as implemented in OpenCV.

We thus obtain a stack of frames in texture space where each UV coordinate indexes a Texel’s reflectance trace over time. Given a Texel at texture coordinate (u, v) , we denote its reflectance trace as the set of direction-radiance pairs observed over time (see fig. 3.5a for a plot of a typical reflectance trace):

$$\mathcal{T}_{u,v} = \{(\vec{\omega}_{bs}, L_{u,v}(\vec{\omega}_{bs}))\} \quad (3.2)$$

where $\vec{\omega}_{bs}$ is the back-scattering direction (see section 3.2.2) and $L_{u,v}(\vec{\omega}_{bs})$ is the observed reflected radiance at the current Texel position and direction.

3.2.2 Back-scattering direction estimation

In our method, we define the 3D position of the mobile device as an orthonormal right-handed coordinate frame, where \vec{x} points to the right of the device, \vec{y} to the top and \vec{z} out of the screen when the device is in landscape mode (see fig. 3.3). The back-scattering direction can then be approximated as the \vec{z} -direction of this coordinate system. In order to estimate this coordinate system, we investigated two complementary methods.

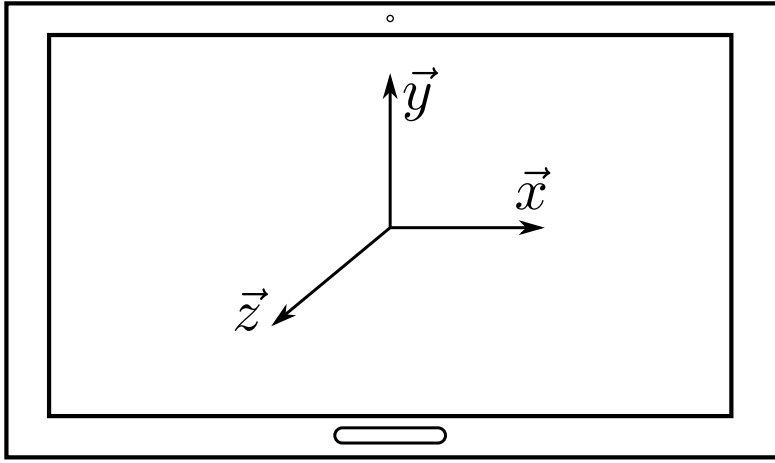


Figure 3.3: **Calibration geometry:** Device-centred coordinate system.

Sensor-based Tracking

Typical mobile devices contain a wide variety of Inertial Measurement Unit (IMU) sensors that can aid in estimating their relative position. Such information is generally used to switch the device’s display orientation between landscape and portrait depending on its orientation or to control games and VR applications by moving the device in space. In our work, we propose to take advantage of such inertial data to estimate the relative orientation of the back camera-flash pair of the device with respect to the sample. The Android Application Programming Interface (API) provides a wealth of functions to leverage the many IMU sensors available on the device.

We start by querying the device’s rotation matrix, which transforms a vector from the device’s coordinate system to the world’s coordinate system, defined as the right-handed orthonormal basis $[\vec{x}_w, \vec{y}_w, \vec{z}_w]$ where:

- $\vec{x}_w = \vec{y}_w \times \vec{z}_w$ is tangential to the ground at the device’s location and roughly points East towards the horizon.
- \vec{y}_w is tangential to the ground at the device’s location and points towards the magnetic North.
- \vec{z}_w points towards the zenith at the device’s location.

The corresponding API call (*getRotationMatrix*) requires inputs from the gravity sensor (which defines \vec{z}_w) and the geomagnetic sensor (which defines \vec{y}_w).

To derive the device's orientation with respect to the sample, we further rely on the *getOrientation* API call which provides the azimuth (angle of rotation around the z-axis), pitch (angle of rotation around the X-axis) and roll (angle of rotation around the y-axis) in the device-centred coordinate system (fig. 3.3) and save the pitch and roll angles for each acquired frame, where pitch == roll == 0° corresponds to a view of the sample at normal incidence. As the user measures different lines of latitude of the upper hemisphere surrounding the sample, we interpret changes in pitch and roll as changes in θ and ϕ to obtain the direction of back-scattering for each acquired frame.

Vision-based Tracking

Alternatively, when the mobile device lacks the required sensors and/or the material sample exhibits sufficient texture, vision-based tracking can be used to estimate the back-scattering direction. Inspired by the recent success of Parallel Tracking And Mapping (PTAM) [78] for unstructured light fields acquisition [79] and for augmented reality with surface light fields [80], we estimate the camera parameters (both intrinsics and extrinsics) using PTAM. While limited to sufficiently textured material samples, vision-based 3D tracking provides the 3D position of the camera (and thus light source) as well as the 2D direction it is pointing towards, as opposed to the sensor-based approach which only provides the latter. This allows to take local lighting effects and camera perspective into account, producing more accurate reflectance estimates in theory. However, in section 3.5, we show visually good qualitative agreement between both approaches.

3.2.3 Radiometric calibration

Radiometric calibration is essential to ensure that the measured intensities are coherent across different views. Given the hand-held nature of our approach, differences in intensity can arise due to changes in distance to the sample, as well as due to changes in ambient lighting. Furthermore, the laws governing reflection are linear while in general, camera sensors are not. It is thus paramount to be able

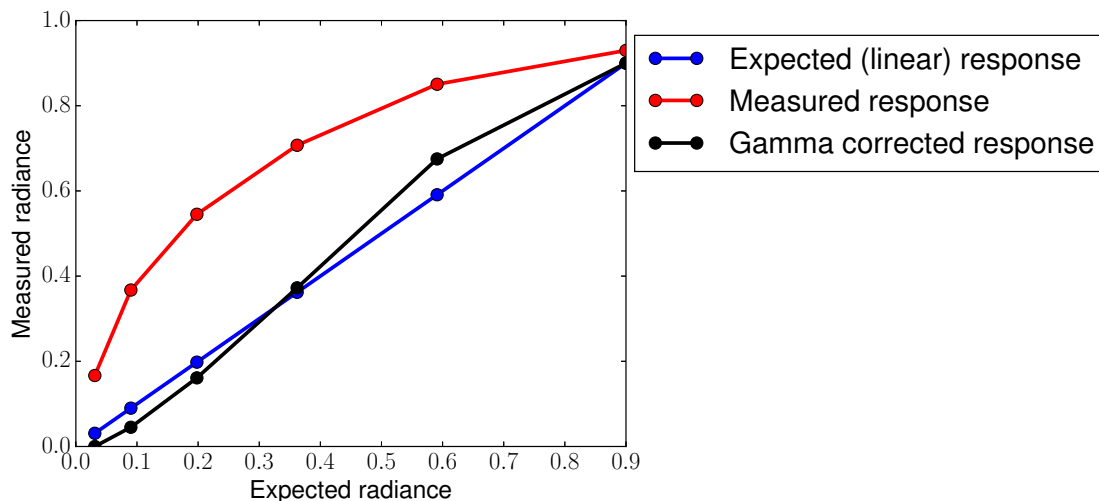


Figure 3.4: **Back camera response:** Given the limited control on exposure of our device, we instead fit a gamma function ($\gamma = 3.2$) to the observed radiance of the X-Rite ColorChecker®’s grayscale gradient.

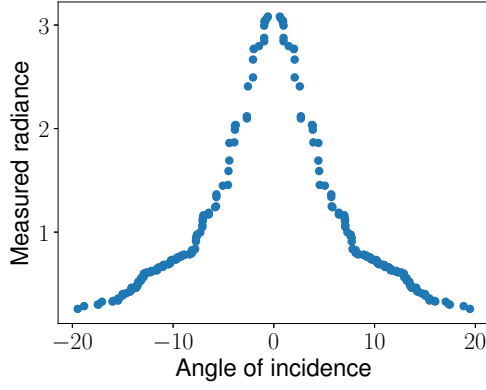
to undo any non linearity caused by the camera’s sensor.

A common method to obtain the response curve of a camera is to capture the same scene at multiple exposures to recover a mapping from measured intensity to true radiance [11, 81]. However, given the limitations in exposure control at the time, we could not follow such an approach. Instead, we rely on the observations of the grayscale gradient of the X-Rite ColorChecker® chart to which we fit a gamma curve to linearise the measured radiance data. As can be seen in fig. 3.4, this calibration is far from perfect, but we found it to be accurate enough for our purpose. Furthermore, since Android 5 “Lollipop”, the API has been extended to allow fine-grained control of the camera hardware, in particular for better camera exposure control.

3.3 Reflectance recovery

3.3.1 Surface normal estimation

The key insight in estimating surface normals with our approach is to observe that both diffuse and specular peak responses align with the surface normal when imaging reflectance in the back-scattering direction. Given the hand-held nature of



(a) Reflectance trace at the Texture Element (Texel) location indicated by a dot in (b)

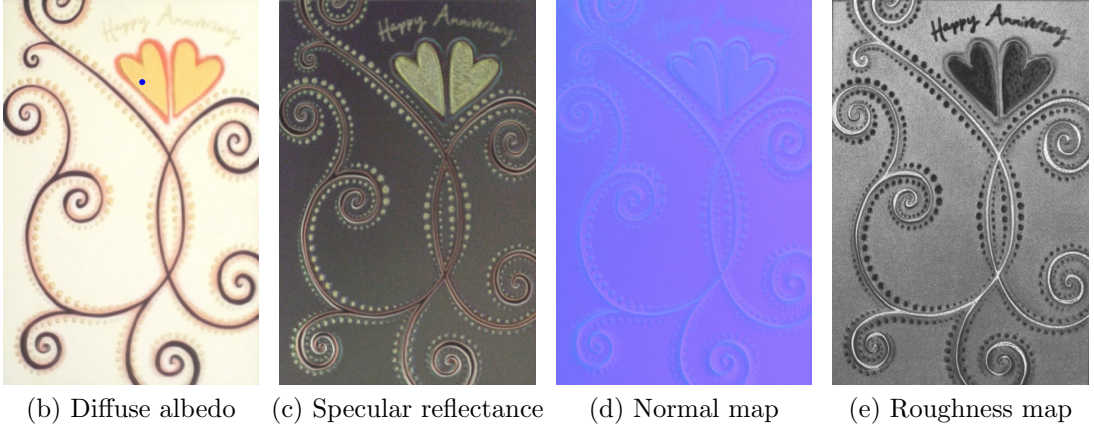


Figure 3.5: **Anniversary greeting card:** Example reflectance maps (bottom) recovered from our mobile free-form acquisition setup by decomposing the measured reflectance trace (a) to recover per-Texel diffuse albedo and specular reflectance (b,c), normals (d) and roughness (e).

our free-form acquisition setup however, sampling the exact direction of maximum reflectance is not guaranteed, but sampling close to that direction is extremely probable for rough specular materials.

We therefore estimate the surface normals as a weighted average of the top 20% brightest observations. If we denote $\mathcal{T}_{u,v}^+$ the trace of the 20% brightest observations of a Texel, we have:

$$\vec{n}_{u,v} = \frac{1}{\sum_i w_i} \sum_{i=0}^{|\mathcal{T}_{u,v}^+|} w_i \vec{\omega}_{bs,i} \quad (3.3)$$

where the weight w_i is the observed radiance at the i^{th} observation in $\mathcal{T}_{u,v}^+$ and $\vec{\omega}_{bs,i}$ the corresponding direction of back-scattering, thus giving more importance to directions closest to the true normal.

3.3.2 Diffuse component estimation

Given the nature of our measurement setup, our data is particularly well suited for fitting to Ashikhmin and Premože’s d-BRDF [75] (see fig. 3.5a). However, unlike them, most of our datasets have a significant diffuse component which first needs to be subtracted from the data before fitting the specular residual to the d-BRDF (eq. (2.24)).

In one of their examples where diffuse reflection was significant, Ashikhmin and Premože proposed a heuristic approach to estimate the diffuse component assuming a Lambertian surface, by averaging the measured back-scattered response in a region outside of the specular peak. In our work, we also follow a heuristic approach but prefer the median to the mean operator.

The key insight in our diffuse albedo estimation is that diffuse reflection is low frequency while specular reflection varies rapidly during data acquisition. Given an observation of a Texel’s trace without the top 20% brightest intensities $\mathcal{T}'_{u,v} = T_{u,v} \setminus T_{u,v}^+$ (which we know to be specular-dominated), we find the observation that corresponds to the median intensity of $\mathcal{T}'_{u,v}$ which we divide by the foreshortening term ($(\vec{n} \cdot \vec{\omega}_{bs})$) to recover the diffuse albedo. The median operator, just as the mean operator, plays the role of a low-pass filter. However, we found the median operator to be more convenient in our approach, as any slight misalignment during the registration phase (section 3.2.1) would cause visible ghosting artefacts with the mean operator.

3.3.3 Specular component estimation

From our estimates of surface normals and diffuse albedo, we compute a per-frame diffuse only reflection component which we subtract from the observed data to obtain a novel stack of registered frames with only the specular response left.

As shown in [75], these measurements are directly proportional to the BRDF and could be used directly as a Look-Up Table (LUT) for rendering. However, such an approach is only viable for homogeneous materials where one LUT is enough

to represent the reflectance of the whole surface.

Instead, to reduce memory foot-print, we start by fitting the specular-only response frames to Ashikhmin and Premože’s d-BRDF (eq. (2.23)) to obtain the surface’s specular BRDF $f_s(\vec{\omega}_{bs})$ in the back-scattering direction and estimate the specular reflectance map $F(0^\circ)$ (see fig. 3.5c) as the hemispherical integral of the observed diffuse-free BRDF. We then further fit a GGX distribution (eq. (2.22)) by non linear least-squares, to obtain the surface’s specular roughness map (see fig. 3.5e).

3.3.4 Results

In the following section, we present some results of spatially-varying planar greeting cards acquired with the presented free-form acquisition setup. We chose those three examples (“new job”, “twenty-one” and “anniversary” greeting cards) as they exhibit a rich spatial variation in their reflectance properties.



Figure 3.6: **Free-form acquisition - maps:** Reflectance maps recovered from our free-form acquisition setup for three spatially-varying rough specular greeting cards.

Figure 3.6 shows reflectance maps of our three test samples acquired with the proposed hand-held, free-form acquisition method presented in the previous sections. As can be seen, the recovered reflectance maps fully capture the rich per-Textel variations of the real materials.

To validate our approach, we compare rendered imagery of our samples under point light illumination to photographs of the sample acquired in the same conditions. Figure 3.7 presents photo-rendering comparison of our test samples for back-scattered reflection at normal incidence. While the results are qualitatively good, a close-up inspection shows blur in the rendered imagery. Looking at the top of the “new job” greeting card for example (fig. 3.7, first row), the photograph depicts the phrase “congratulations on your” in white, with crisp details while the same phrase is significantly blurred in the corresponding rendering. This is a direct consequence of the hand-held nature of the setup: any motion blur and

slight misalignment after the registration phase (section 3.2.1) results in blur in the recovered reflectance maps. To overcome this limitation, we thus propose a post-processing step to enhance the recovered reflectance maps (section 3.4), which helps improve the impression of realism in rendering.




Figure 3.7: **Free-form acquisition - results:** Photo-rendering comparison under point light illumination at normal incidence (back-scattered reflection).

3.4 Surface detail augmentation

To account for the loss of details in the renderings from our acquired reflectance maps, we follow an approach akin to that of Beeler et al. for mesoscopic augmentation of facial geometry [82]. Unlike in their approach, we not only augment the recovered normals but also the diffuse and specular reflectance maps. Similarly to the approach of Beeler et al., the enhanced details (see section 3.4.1 for our algorithm) are not exact but provide plausible details for photo-realistic rendering applications. Section 3.4.2 presents comparison results before and after applying the enhancement step for some of our samples.

3.4.1 Algorithm

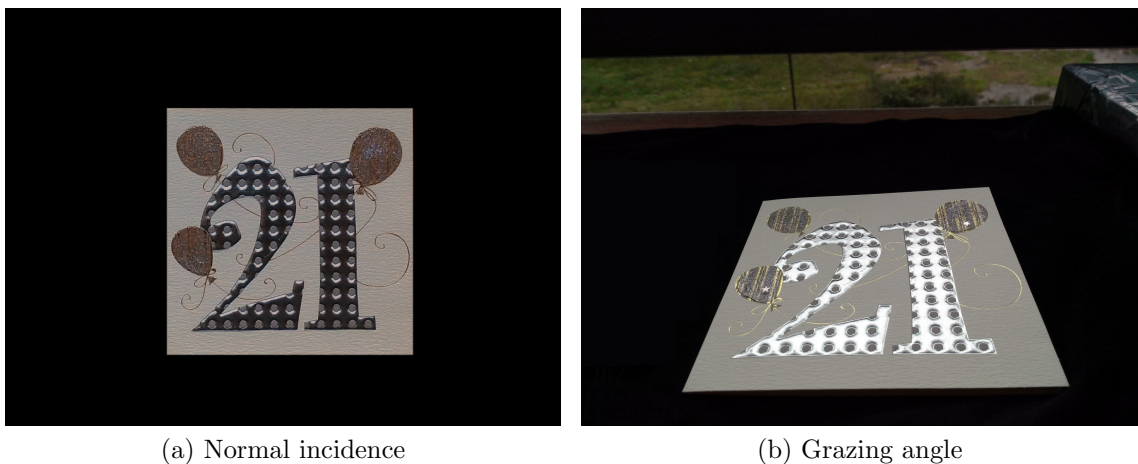


Figure 3.8: **Surface details enhancement input:** We take two additional close-up observations of the sample under natural illumination through a window to recover additional mesoscopic surface details not visible in the reflectance maps recovered from free-form acquisition.

In order to recover the missing high-frequency details from our reflectance maps, we take two additional close-up views of the sample under natural illumination through a window with our device’s high-resolution back-facing camera: one at normal incidence to excite mostly the diffuse component of reflection (fig. 3.8a) and one at grazing angle facing the window (fig. 3.8b) to excite specular reflectance.




Figure 3.9: **Surface detail enhancement algorithm:** We add mesoscale details (middle column) to the maps recovered from free-form acquisition (first column) to obtain highly detailed reflectance maps (third column). Zoomed-in crop to highlight details.

After registration to the normal incident view, we employ these additional observations to extract details at the mesoscopic scale by subtracting from each high-resolution images, a copy of itself blurred by a Gaussian kernel. In our experiments, we set the Gaussian kernel’s width to 4, but this value could be adapted to extract more or less details. This step acts as a high-pass filter, preserving only high frequency details (fig. 3.9: column 2, rows 1 and 2). We then add the high-pass filtered enhancement map taken at normal incidence (resp. grazing angle) to the diffuse (resp. specular) reflectance maps, to transfer back the missing high

frequency details (fig. 3.9: column 3, rows 1 and 2). The added benefit of this step is that the maps get up-scaled to the resolution of the high resolution view at normal incidence.

Finally, our grazing angle observation is further employed to enhance the surface normals. We compute its gradients along the \vec{x} and \vec{y} axes of the image plane (fig. 3.9: column 2, row 3) and add them to \vec{n}_x and \vec{n}_y (the components of the surface normal along the \vec{x} (resp. \vec{y}) direction) respectively. After normalisation (i.e. ensuring $\vec{n}_x^2 + \vec{n}_y^2 + \vec{n}_z^2 = 1$), we obtain enhanced surface normals (fig. 3.9: column 3, row 3) which better reproduce the fine-scale details of the real material.

3.4.2 Results

While not exact, our detail enhancement step helps in conveying a better impression of realism when rendering the acquired material. Figure 3.10 presents a photo-rendering comparison before (2nd column) and after (3rd column) surface detail enhancement.




Figure 3.10: **Surface detail enhancement results:** Here we compare photos of our samples (a) to renderings before (b) and after (c) surface detail enhancement. Zoomed in crop to highlight details.

As can be seen, both renderings capture the overall reflectance properties of the material but a closer inspection at the third column demonstrates sharper details that are closer to the corresponding photographs and convey a better sense of realism. The phrase “congratulation on your” on the “new job” sample (first row), after detail enhancement, now appears in full details. The “twenty-one” sample (second row) with enhancement better captures small details such as the shiny swirls and creases on the base paper.

3.5 Discussion and limitations

In section 3.2.2, we discussed two methods for estimating the back-scattering direction using the device’s internal sensors or a vision-based tracking method [78]. While the latter approach provides more accurate tracking and allows to take into account local lighting effects, we found that in practice, the sensor-based tracking is already quite accurate for rendering applications.

Figure 3.11 provides a photo-rendering comparison for the “anniversary” greeting card for maps estimated with sensor-based tracking (fig. 3.11f) and vision-based tracking using PTAM (fig. 3.11d) and shows good qualitative agreement in both cases. Here, we also demonstrate good photo-rendering matching for novel views (second and third rows).

In our experiments, we found our free-form flash-based approach to work well in practice for rough specular samples. As the material’s specular response gets sharper however, the sampling rate imposed by the BRDF’s high frequency becomes impractical to match with a hand-held device. To circumvent this limitation, we thus investigated how the device’s LCD panel could be used as an extended source of illumination, which will be the main focus of chapter 4. Our method is inspired by the LCD-based reflectometry approach presented of Ghosh et al. [5].



Figure 3.11: **Mobile Surface Reflectometry - validation:** Comparison of renderings with normal maps obtained from sensor-based tracking (c) and 3D tracking (d), to a photograph (b). Here, we also demonstrate good photo-rendering matching for novel view renderings ((b)-(f)) in the mirror direction (a).

Chapter 4

Mobile Surface Reflectometry - LCD-based Acquisition

The previous free-form, flash-based reflectometry method (chapter 3) samples the surface reflectance from a set of discrete directions. As with other sampling-based methods, this places a limit on the sharpness of the specular surface reflectance that can be accurately recovered. In this chapter, we propose an alternative mobile reflectometry solution for sharp specular materials that utilises the mobile device’s LCD panel as a source of extended illumination, to project spherical gradient illumination patterns (section 4.1). However, due to the small size of the LCD panel and the off-centre location of the front camera, only a small 5x5 cm section of the material sample can be directly recovered (section 4.2). We thus rely on a novel appearance transfer method to extend the acquired surface reflectance to larger material samples (section 4.3)

4.1 Data acquisition

Similarly to our free-form acquisition setup, we employed a Fujitsu Stylistic M532 10” Android tablet for our LCD-based measurement setup (fig. 4.1b) with a 2 MP front facing camera. Similar to the work of Ghosh et al. [5], we project gradient illumination patterns (fig. 4.1b) from the device’s display and take advantage of

the inherent linear polarisation of the device’s LCD screen for diffuse-specular separation.

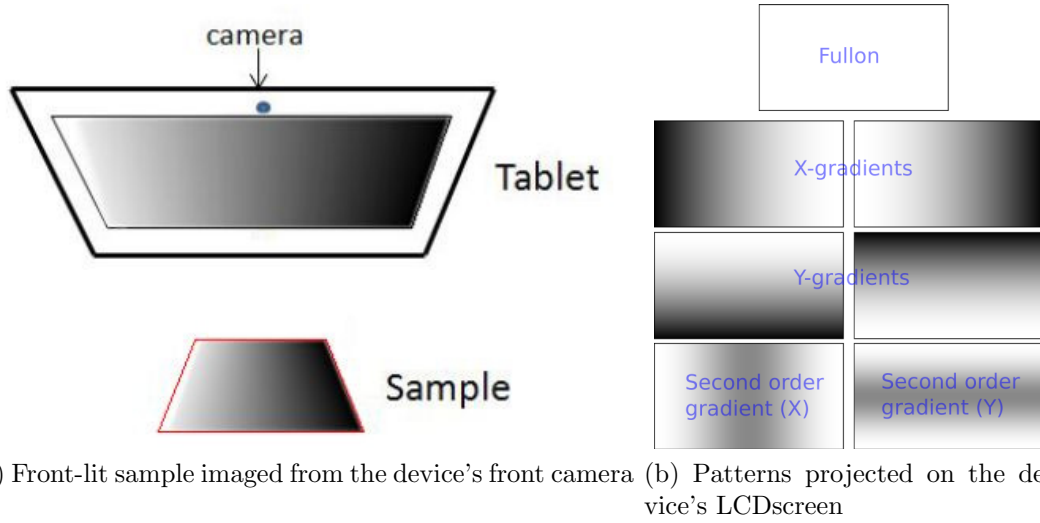


Figure 4.1: **Mobile LCD-based acquisition:** We project gradient illumination patterns (b) from the device’s screen and image the sample from the front facing camera (a).

We statically mount our tablet with its front camera facing down toward the sample at a height of roughly 45 cm and proceed to capture a first set of pictures under the different lighting patterns shown in fig. 4.1b with a sheet of linear polarising filter mounted in front of the camera so as to be cross-polarised with respect to the device’s screen polarisation. In this configuration, we are imaging only the light that’s being reflected diffusely. In a second time, we rotate the polariser in front of the camera by 90° so as to image the mixed diffuse and specular signal reflected by the sample under each lighting pattern. For our device, we found the polarisation axis of the screen to be at 45° with respect to the $\vec{x} - \vec{y}$ plane shown in fig. 3.3. This is to be expected for any mobile device (phone, tablet) with a linearly polarised LCD as they are likely to be used outdoors in different orientations, where the user might be wearing vertically polarised sunglasses (see section 5.1.2 for an explanation). Therefore, vertically or horizontally polarising a mobile device’s screen would make it unusable with polarised sunglasses in one of portrait or landscape mode.

Essentially, each lighting pattern allows us to image the 0^{th} , 1^{st} and 2^{nd}




Figure 4.2: **Mobile LCD-based data:** We acquire two sets of data under gradient illumination, taking advantage of the inherent polarisation of the device’s screen for diffuse-specular separation.

order spherical moments of the material’s reflectance function. Similarly to Ghosh et al. [5], we denote these spherical moments by L_0 , L_1 and L_2 respectively. The rationale for measuring with those particular lighting patterns is that the 0^{th} and 1^{st} moments correspond to the albedo and surface normals of the observed material [57], while the 2^{nd} order moments are related to the specular roughness [5] (i.e. the width of the specular lobe). Figure 4.2 shows the cross-polarised (first row) and parallel-polarised (second row) data acquired for a shiny bronze medal under gradient illumination. From these, we can infer reflectance maps of the material for rendering (section 4.2.1). Note that we also employ the observation of the X-Rite ColorChecker®’s grayscale gradient under full-on illumination for radiometric calibration, as in section 3.2.

4.2 Processing pipeline

In this section, we first cover the implementation details necessary to extract reflectance maps (section 4.2.1) from data such as presented in fig. 4.2. For an in-depth derivation of the underlying theory, we refer the reader to the original paper from Ghosh et al. [5] for which we show a successful application using a self-contained mobile acquisition setup. We then present reflectance maps inferred from the acquired data (section 4.2.2) and show how our detail augmentation step from section 3.4 can be applied to overcome the lack of resolution of the front

facing camera (2MP for our device). Finally, we discuss limitations inherent to the acquisition setup in terms of maximum measurable dimensions of the material under study and introduce a novel appearance transfer framework that is able to infer plausible reflectance properties for larger samples that cannot be measured directly. We do so by combining observations of a small exemplar crop of the material with observations of the sample under natural illumination through a window, similar to our detail augmentation step presented in section 3.4.

4.2.1 Reflectance recovery

Diffuse-specular separation

In our setup, diffuse-specular separation comes “for free”, as we are imaging the sample under both cross-polarised light, to cut off any specular reflection and parallel-polarised light to image a mix of diffuse and specular reflectance. The diffuse and specular reflectance maps are thus obtained from the 0^{th} order moment images (fig. 4.2a) after calibration, as:

$$\begin{aligned}\rho_d &= L_0^\times \\ F(0^\circ) &= L_0^\parallel - \rho_d\end{aligned}\tag{4.1}$$

where L_0^\times and L_0^\parallel are the 0^{th} order moments observed under cross-polarised and parallel-polarised illumination respectively.

Surface normal estimation

From our data, we could, in principle, compute diffuse and specular normals from the eight 1^{st} order moment images (Figures 4.2b to 4.2e) and render our materials with hybrid normals similar to Ma et al. [57]. However, in our experiments, we found the diffuse normals (fig. 4.3a) to be too noisy, due to limited levels of illumination allowed by the tablet’s screen as well as the limited extent of the illumination which does not cover the full hemisphere (Ω^+) surrounding the surface. Furthermore, as pointed out by Ma et al., specular normals (fig. 4.3b) reflect the




Figure 4.3: **Hybrid normals problem:** Due to the low intensity of light emitted from the device’s LCD screen and limited extent of the screen over the sample’s upper hemisphere (Ω^+), the diffuse normals (a) are too noisy for rendering. Instead, we render both diffuse and specular reflections with specular normals (b).

true shape of the material. For these reasons, we only relied on specular normals (hereafter referred to as simply normals) for both types of reflections in rendering.

To obtain our surface normal maps, we first subtract each cross-polarised L_1 images from its corresponding parallel-polarised L_1 counterpart to obtain specular only 1st order moment frames:

$$L_{1,\bullet}^s = L_{1,\bullet}^{\parallel} - L_{1,\bullet}^{\times} \quad (4.2)$$

where \bullet can be any of \vec{x} , $-\vec{x}$, \vec{y} , $-\vec{y}$. From eq. (4.2), we separately compute the Cartesian components of each per-Texel surface normal $\vec{n} = \begin{bmatrix} n_x & n_y & n_z \end{bmatrix}^T$ as follows:

$$\begin{aligned} n_x &= \frac{L_{1,+\vec{x}}^s - L_{1,-\vec{x}}^s}{F(0^\circ)} \\ n_y &= \frac{L_{1,+\vec{y}}^s - L_{1,-\vec{y}}^s}{F(0^\circ)} \\ n_z &= \sqrt{1 - (n_x^2 + n_y^2)} \end{aligned} \quad (4.3)$$

Specular roughness

Finally, we obtain specular roughness from the 2nd order moment images (figs. 4.2f and 4.2g) by first subtracting each cross-polarised L_2 images from its corresponding

parallel-polarised L_2 counterpart to obtain $L_{2,\bullet}^s$, and compute roughness estimates along the orthogonal \vec{x} and \vec{y} directions:

$$\begin{aligned}\sigma_x^2 &= \frac{L_{2,\vec{x}}^s}{F(0^\circ)} - \left(\frac{L_{1,\text{sign}(n_x)\vec{x}}^s}{F(0^\circ)} \right)^2 \\ \sigma_y^2 &= \frac{L_{2,\vec{y}}^s}{F(0^\circ)} - \left(\frac{L_{1,\text{sign}(n_y)\vec{y}}^s}{F(0^\circ)} \right)^2\end{aligned}\tag{4.4}$$

The isotropic specular roughness is then obtained as the magnitude of the 2D-vector $\begin{bmatrix} \sigma_x^2 & \sigma_y^2 \end{bmatrix}^T$.

4.2.2 Results




Figure 4.4: **LCD-based mobile reflectometry - direct capture:** Reflectance maps acquired for small shiny objects as per the protocol outlined in section 4.1. The red line delimits the frontier between reflectance maps without post-processing (left) and reflectance maps obtained after surface detail augmentation as presented in section 3.4 (right).

Figure 4.4 presents reflectance maps recovered from the protocols outlined in sections 4.1 and 4.2 for two highly specular samples which could not be acquired with the method presented in chapter 3. As can be seen, the resulting reflectance maps present good qualitative spatial variations but suffer from a slight blur of the high frequency details, due to the generally lower resolution of the front camera on mobile devices (2MP for our device).

To overcome this limitation, we thus employ the detail enhancement step presented in section 3.4 to “up-res” the acquired reflectance maps and recover the high frequency details. The effects of the detail augmentation are especially noticeable in the normal map of the “Intel coin” (fig. 4.4c, 2nd row) which looks almost flat with direct capture and on the background of the “Bronze medal” (Figures 4.4c to 4.4b, 1st row) where the small bumps are recovered by the enhanced maps.

From these maps, we are able to render photo-realistic virtual clones of the real object under novel illumination. Figures 4.5 and 4.6 show the enhanced maps (1st row) obtained for both the “Bronze medal” and “Intel coin” samples and their photo-rendering comparison (2nd row). While both renderings produced are qualitatively close to their corresponding photographs, with matching highlights, we notice a slight mismatch in colour tone for both samples which we attribute to the low quality of colour rendition from the front facing camera under low lighting conditions (the lighting comes mainly from the device’s screen).




Figure 4.5: **Bronze medal:** Reflectance maps ((a)-(d)) recovered from our mobile LCD-based acquisition setup. The maps allow good qualitative renderings (f) with highlights matching real photographs (e).





Figure 4.6: **Intel coin:** Reflectance maps ((a)-(d)) recovered from our mobile LCD-based acquisition setup. The maps allow good qualitative renderings (f) with highlights matching real photographs (e).

4.3 Appearance transfer



(b) Example of a sample too large for direct capture



(c) Photograph under point light source showing low-rank variations in specular response

Figure 4.7: **Limited measurable size:** The non ideal overlap between view and light frustum (a) makes it impossible to directly measure larger samples (b). Our appearance transfer approach is directly motivated by the observation that spatially-varying specular appearance tends to be low-rank (c).

The pipeline presented in section 4.2 is, in practice, limited to small objects where the maximum allowed size of the object being scanned depends on the size of the device used for measurements; in theory, this maximum limit corresponds roughly to objects of the size of the LCD panel used as a source of illumination. However, with our mobile device, we found this limit to be much lower at around 5x5 cm for our 10" tablet. This is a problem for any smartphone/tablet and stems from the limited overlap between the front camera's view frustum and extent of the screen's illumination (fig. 4.7a). The reason for this is that the front cameras on such devices are meant to be focused on the user's face for video chats, selfies, etc. Therefore, manufacturers tend to build mobile devices such that their front camera's optical axis points upwards.

For this reason, our proposed setup is not able to directly measure large samples such as the “cha” tea box cover presented in fig. 4.7b. However, we notice that spatially-varying materials generally exhibit rich variations in diffuse texture while their variations in specular appearance is generally low rank, as can be seen in fig. 4.7c, a photograph of the tea box cover under flash illumination. For this sample, the specular reflectance properties can be classified in two clusters: the text and the metal base.

From those observations, we propose a novel appearance transfer mechanism (section 4.3.1) based on observations of the sample under natural illumination to transfer the specular properties of a small representative crop acquired as per the protocol described in section 4.2. Note that we only transfer specular reflectance and roughness properties and obtain diffuse and normal maps from our free-form acquisition setup (chapter 3).

4.3.1 Algorithm

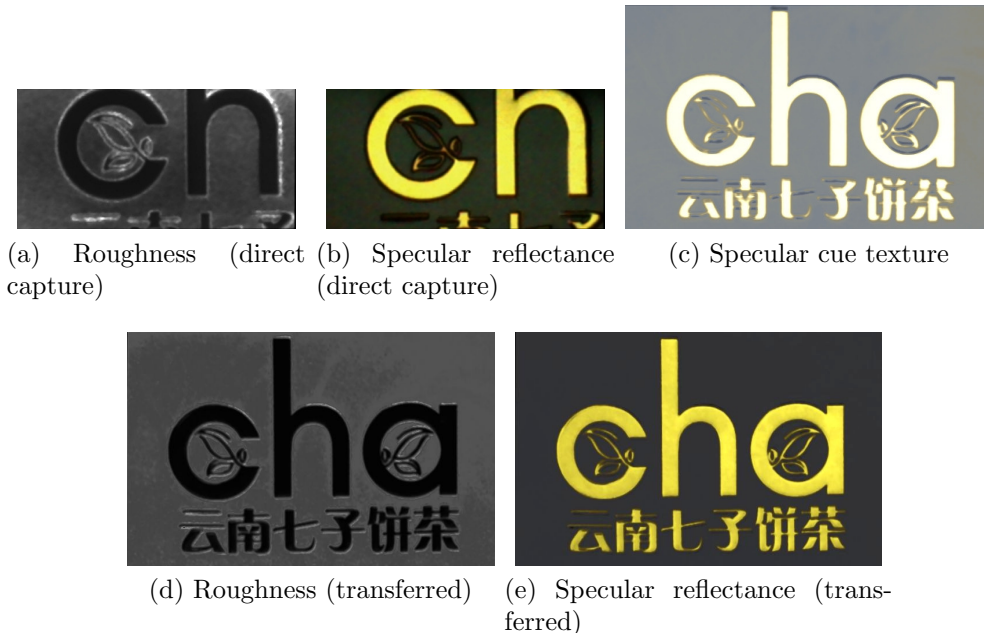


Figure 4.8: **Appearance transfer:** From a small crop obtained by direct capture ((a),(b)) and a specular cue texture (c) we transfer both roughness (d) and specular reflectance (e) properties to a larger portion of the sample.

Recall, from section 3.4, we added two additional observations at normal inci-

dence and grazing angle under natural illumination through a window (fig. 3.8) to enhance the appearance of diffuse and specular reflectance maps. The insight then was that the normal incidence observation was mainly imaging diffuse reflection while the grazing angle photo was dominantly specular.

For our appearance transfer mechanism, we again take advantage of such observations: in addition to the specular reflectance ($F(0^\circ)^{direct}$) and roughness maps (σ^{direct}) (figs. 4.8a and 4.8b) measured for a small crop of the material by the protocol outlined in sections 4.1 and 4.2, we add two additional photos under natural illumination at normal and grazing incidence. Subtracting the latter from the normal incidence photo provides a specular cue texture (fig. 4.8c, I^{cue}) which drives the appearance transfer described below in pseudo-code:

Algorithm 1 Appearance transfer

```

1: procedure REGISTER(Input, Target)
2:   Register Input to Target
3:   return  $Input_{reg}$ 
4: procedure COMPUTEBBOX(Input)
5:   Compute BBox from Input
6:   return BBox
7: procedure CROP(Input, BBox)
8:   Crop Input within BBox
9:   return  $Input_{cropped}$ 
10: procedure COMPUTELUT(Input, LUTSize)
11:   Down sample and vectorise Input to match LUTSize
12:   return LUT
13: procedure PRE-PROCESS
14:    $F(0^\circ)_{reg} \leftarrow$  REGISTER( $F(0^\circ)^{direct}$ ,  $I^{cue}$ )
15:    $\sigma_{reg} \leftarrow$  REGISTER( $\sigma^{direct}$ ,  $I^{cue}$ )
16:    $BBox \leftarrow$  COMPUTEBBOX( $\sigma_{reg}$ )
17:    $\sigma_{LUT} \leftarrow$  COMPUTELUT( CROP( $\sigma_{reg}$ , BBox), 4096)
18:    $F(0^\circ)_{LUT} \leftarrow$  COMPUTELUT( CROP( $F(0^\circ)_{reg}$ , BBox), 4096)
19:    $I_{LUT}^{cue} \leftarrow$  COMPUTELUT( CROP( $I^{cue}$ , BBox), 4096)
20:   return  $\sigma_{LUT}$ ,  $F(0^\circ)_{LUT}$ ,  $I_{LUT}^{cue}$ ,  $I^{cue}$ 
21: procedure MAIN
22:    $\sigma_{LUT}$ ,  $F(0^\circ)_{LUT}$ ,  $I_{LUT}^{cue}$ ,  $I^{cue} \leftarrow$  PRE-PROCESS()
23:   Initialise  $F(0^\circ)^{transferred}$  and  $\sigma^{transferred}$  to be the same size as  $I^{cue}$ 
24:   for  $x \in [0, width(I^{cue})]$  do
25:     for  $y \in [0, height(I^{cue})]$  do
26:       Find index  $i$  of nearest neighbour for  $I^{cue}(x, y)$  in  $I_{LUT}^{cue}$ 
27:        $F(0^\circ)^{transferred}(x, y) \leftarrow F(0^\circ)_{LUT}[i]$ 
28:        $\sigma^{transferred}(x, y) \leftarrow \sigma_{LUT}[i]$ 

```

We start by registering $F(0^\circ)^{direct}$ and σ^{direct} to the specular cue texture and compute a LUT for each by mapping the observed intensity in the specular cue texture to its corresponding Texel in the cropped specular reflectance and roughness maps directly measured. For each Texel of the cue texture that does not have a mapped value, we find its nearest neighbour in intensity and assign the corresponding value to each of $F(0^\circ)^{transferred}$ and $\sigma^{transferred}$. Note that in our experiments, we found a LUT size of 4096 to give the best trade-off between computation time, noise reduction and overall detail preservation.

4.3.2 Results

Figure 4.9 presents reflectance maps recovered for two samples larger than the size of the measurement frustum. The normal maps (fig. 4.9c) and diffuse albedo maps (fig. 4.9a) are obtained from the free-form acquisition method presented in chapter 3, while the specular reflectance (fig. 4.9b) and roughness (fig. 4.9b) maps are transferred from the crops shown in the last row of fig. 4.9.



Figure 4.9: **Appearance transfer - maps:** From observations of a small representative patch of the material (e-h), we transfer specular reflectance (c) and roughness (d) properties to the entire sample as per the protocol described in section 4.3.1.

We validate our appearance transfer approach by comparing renderings (fig. 4.10b) obtained from the reflectance maps presented in fig. 4.9 to photographs (fig. 4.10a) and show good visual agreement, despite the limitations of the acquisition protocol. Furthermore, our reflectance maps generalise well to novel view/lighting conditions (fig. 4.10c), producing plausible rendered imagery.



Figure 4.10: **Appearance transfer - results:** Despite the simplicity of the method, our recovered reflectance maps produce compelling renderings (b,c) that well match real photographs (a).

Mobile Surface Reflectometry - Summary

In part II, we introduced two novel approaches for mobile surface reflectometry. We proposed a free-form, hand-held solution (chapter 3) that exploits the near co-axial configuration of camera and flash on a typical mobile device, suited for rough specular materials. We further presented a method to enhance the acquired reflectance maps from observations of the sample under natural illumination, able to emboss high frequency details, to help convey a better impression of realism.

For highly specular materials, we proposed to employ extended illumination from the device's LCD screen, projecting spherical gradient illumination patterns, in conjunction with the front camera to acquire surface reflectance. For diffuse-specular separation, we exploited the inherent polarisation of the device's screen. We further proposed a novel appearance transfer method that combines controlled measurements of an exemplar section of the entire object with observations of the sample under natural lighting.


We demonstrated that the proposed mobile surface reflectometry solutions achieve high quality reflectance and mesostructure reconstructions on a wide range of planar material samples. Those methods are however limited to indoor environments where the ambient lighting can be controlled not to exceed 40 LUX. This limits the use of such technology outdoors where the illumination from the sun and sky is dominant. The next section of this thesis (part III) is set out to alleviate those limitations to allow the measurements of the reflectance properties of permanent on-site structure under outdoors natural illumination.

PART III:


OUTDOORS REFLECTOMETRY [2]



(a) Specular drain cover



(b) Stone pavement



(c) Red book cover



(d) Sketch book cover

Figure II: **Polarisation imaging reflectometry in-the-wild:** High-resolution renderings from reflectance and normal maps of planar samples obtained by exploiting multiple polarisation observations under uncontrolled outdoor illumination.

Outdoors Reflectometry - Prelude

Both of the mobile surface reflectometry approaches presented in part II require that the device's illumination (flash or LCD screen) dominate the ambient light levels, which can be achieved in a generally dimly lit indoor environment, thus relaxing the dark room requirements of previous approaches to some extent. However, the methods are still limited to rather low amount of ambient light (15 to 40 LUX) which tells us that active illumination approaches are not, in general, suitable for reflectometry in uncontrolled environments such as natural illumination outdoors.

In the remainder of this thesis, we thus set out to investigate a method for passive reflectometry in general outdoors conditions. In particular, we exploit observations of the state of polarisation of light reflected under natural illumination to derive high resolution reflectance maps of permanent on-site samples that would otherwise be impossible to scan with existing methods. Our choice of polarisation imaging as a tool for reflectometry is motivated by two key observations:

- Open sky shows strong patterns of linear polarisation, due to single scattering of photons from the sun by molecules (nitrogen and oxygen mostly) smaller than the wavelength of light. This phenomenon is named Rayleigh scattering, after the British physicist John William Strutt, 3rd Baron Rayleigh, who first observed this phenomenon in the 1800s [83, 84]. Rayleigh scattering explains the blue colour of the sky.
- When overcast, the illumination incident from the sky tends to be depolarised due to multiple scattering events from water molecules. Upon reflection however, light becomes partially linearly polarised due to Fresnel effects,

with maximum polarisation occurring at a specific angle called Brewster angle.

The rest of this segment is organised as follows. We first provide an overview of the mathematics of polarisation in terms of Mueller calculus (section 5.1) and derive the theory for polarisation imaging under incident partial linear polarisation, which we show to be a generalisation of the commonly assumed unpolarised world theory (section 5.2). Finally, we derive a measurement protocol for on-site acquisition of surface reflectance in outdoors environments based on the presented theory (chapter 6). To the best of our knowledge, our method is the first to successfully extract a complete set of reflectance parameters with passive capture in completely uncontrolled outdoor settings.

Chapter 5

Polarisation In Graphics And Vision

Polarisation imaging has been widely studied in the past assuming unpolarised incident illumination for material classification [63], shape estimation [67, 85, 71, 72, 73] and reflectometry [66]. The common guiding thread of all those methods is to measure the polarisation induced by reflection off the material's surface, by taking multiple shots (at least 3) of the object through a rotating linear polariser. In this chapter, we will first cover some background on polarisation from specular reflection and the necessary mathematical tools to deal with polarised light (section 5.1). Note that, in the remainder of this dissertation, we follow the common assumption of depolarising diffuse reflection.

5.1 Background

Polarisation refers to the intrinsic property of a wave to oscillate at a preferred orientation, orthogonal to the direction of propagation. For light waves, polarisation refers to the oscillation of the electric field (\vec{E} -vector) along the direction of propagation $\vec{\omega}$. When at any point in time, no particular direction of oscillation can be determined, the ray of light is said to be unpolarised. Light bulbs are an example of light sources emitting unpolarised light.

One way to obtain polarised light from an unpolarised beam of light is by reflection from a dielectric surface: according to Fresnel equations, a portion of the light gets reflected perpendicular to the plane of incidence (plane containing the normal to the surface \vec{n} and ray of light incident at $\vec{\omega}_i$) while another is reflected parallel to the plane of incidence (fig. 2.3b). At a particular angle of incidence (called Brewster angle), the component of light parallel to the plane of incidence gets completely transmitted, thus reflecting purely linearly polarised light perpendicular to the plane of incidence. At any other angle of incidence, the reflected ray is partially polarised, i.e a mixture of unpolarised and linearly polarised light where the Degree of Polarisation (DOP) depends on the angle of incidence (fig. 2.5, [blue curve](#)). A similar, weaker effect can also be observed upon reflection on metals, where the main difference is that the p-polarised component never goes to 0 and hence complete polarisation is never attained.

In optics, there exist two common mathematical frameworks to describe the polarisation state of light and its changes through polarising elements: Jones calculus, developed by Robert Clarke Jones in 1941 and Mueller calculus, developed by Hans Mueller in 1943 to model the interactions of Stokes vectors with polarising optical elements (section 5.1.1). In this work, as in related literature, we consider the latter framework, as Jones calculus can only account for fully polarised light: Mueller calculus can be seen as a generalisation of Jones calculus to partial polarisation, which covers unpolarised, partially polarised and fully polarised light.

5.1.1 Mueller calculus

Mueller calculus is a mathematical framework used to manipulate Stokes vectors, where the effects of a particular optical element are represented by a 4x4 real-valued matrix called Mueller matrix. The Stokes parameters, developed by George Gabriel Stokes in 1852, are a set values that describe the polarisation of light in terms of its total intensity ($L(\vec{\omega})$), DOP (\mathcal{P}) and the parameters of the polarisation

ellipse (fig. 5.1a) as a 4-vector [86]:

$$\vec{s} = \begin{bmatrix} s_0 \\ s_1 \\ s_2 \\ s_3 \end{bmatrix} = \begin{bmatrix} L(\vec{\omega}) \\ L(\vec{\omega})\mathcal{P} \cos 2\psi \cos 2\chi \\ L(\vec{\omega})\mathcal{P} \sin 2\psi \cos 2\chi \\ L(\vec{\omega})\mathcal{P} \sin 2\chi \end{bmatrix} \quad (5.1)$$

where s_0 is the total intensity of the incident beam, s_1 and s_2 respectively the intensity of 0° and $+45^\circ$ linear polarisation and s_3 the intensity of right circular polarisation. It is common to look at Stokes vectors in a normalised space (i.e. $s_0 = 1$) by dividing each component by s_0 . The Poincaré sphere (fig. 5.1b) then serves as a convenient visualisation tool where the last three normalised Stokes parameters are parametrised in spherical coordinates. In this parametrisation, the Stokes vectors form an orthonormal basis spanning the space of unpolarised ($\mathcal{P} = 0$), partially polarised ($0 \leq \mathcal{P} \leq 1$) and fully polarised ($\mathcal{P} = 1$) light.




Figure 5.1: **Visualising polarisation:** When looking down the propagation direction of a light wave, the tip of its electric field traces an ellipse as it oscillates (a). Stokes parameters are related to the parameters of the polarisation ellipse as per eq. (5.1). They span a 3D space represented in spherical coordinate on the Poincaré sphere (b).

When a beam of light interacts with an optical element, the resulting polarisation state is given by a simple matrix-vector multiplication in Mueller calculus:

$$\vec{s}_o = M\vec{s}_i \quad (5.2)$$

where \vec{s}_i is the polarisation state of the beam of light incident upon the optical element, M encapsulates the optical properties of the polarising element and \vec{s}_o is the resulting Stokes vector.

The combined effect of light interacting with multiple optical elements is simply modelled by stacking the Mueller matrices of each element by the right multiplication rule. Given N polarising elements, each associated with a Mueller matrix $M_k, \forall k \in [1, N]$, their combined effect on the input Stokes vector \vec{s}_i is defined as:

$$\vec{s}_o = M_N M_{N-1} \dots M_k \dots M_1 \vec{s}_i \quad (5.3)$$

For the purpose of this thesis, we will be interested particularly in two types of optical elements: polarising filters (section 5.1.2) and reflectors (section 5.1.3) which are the building blocks of the related work presented in section 2.3 as well as our own method presented in chapter 6.

5.1.2 Polarising filters

Polarising filters are optical elements designed to selectively transmit only light in a particular polarisation state. They exist in two flavours: circular and linear polarisers.

Linear polarisers

Linear polarisers are extremely popular with outdoors hobbyists, especially for water-sports, as they are well suited to reduce glare caused by the sun reflecting at the surface of water: the unpolarised light from the sun becomes largely horizontally polarised upon reflection (due to Fresnel effects). Polarised sunglasses are therefore vertically polarised in order to kill off the unpleasant glare. In graphics and vision, linear polarisation is often used for diffuse-specular separation, where linear polarisers are affixed to both lights and cameras.




Figure 5.2: **Linear polariser:** Geometry of a general linear polariser rotated at an angle ϕ_o from its local coordinate system.

In Mueller formalism, a linear polariser is defined with respect to a local Cartesian coordinate $(\vec{x}, \vec{y}, \vec{z})$ where \vec{x} is the horizontal direction, \vec{y} the vertical direction and \vec{z} the direction of transmission (fig. 5.2). A linear polariser with its optical axis rotated at an angle ϕ_o to the horizon is thus defined as:

$$\begin{aligned}
 M_{pol}(\phi_o) &= M_{rot}(-\phi_o)M_{horiz_pol}M_{rot}(\phi_o) \\
 &= \frac{1}{2} \begin{bmatrix} 1 & \cos 2\phi_o & \sin 2\phi_o & 0 \\ \cos 2\phi_o & \cos^2 2\phi_o & \cos 2\phi_o \sin 2\phi_o & 0 \\ \sin 2\phi_o & \cos 2\phi_o \sin 2\phi_o & \sin^2 2\phi_o & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (5.4)
 \end{aligned}$$

where M_{horiz_pol} is the Mueller matrix for an ideal horizontal polariser:

$$M_{horiz_pol} = \frac{1}{2} \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (5.5)$$

and $M_{rot}(\alpha)$ is the Mueller matrix of a rotator that transforms local into

global coordinates:

$$M_{rot}(\alpha) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos 2\alpha & -\sin 2\alpha & 0 \\ 0 & \sin 2\alpha & \cos 2\alpha & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (5.6)$$

Circular polarisers




Figure 5.3: **Circular polariser:** Geometry of general circular polarisers.

Unlike their linear counterpart, the optical properties of circular polarisers are rotationally invariant. We distinguish two types of circular polarisers, left (fig. 5.3a) and right (fig. 5.3b) circular polarisers, where left or right polarisation indicates the apparent motion of the polarisation ellipse over time.

In the cinematographic industry, circular polarisers are used for passive 3D stereoscopic displays with polarised 3D glasses: each eye sees through a circular polariser with each a different chirality. The projection system then projects each pair of stereoscopic images with the correct chirality to create the impression of 3D. Another common use for circular polarisers is to reduce glare on LCD panels: the unpolarised light incident on the panel passes through a circular polariser and is reflected on the metal base of the screen where its chirality is changed such that it cannot transmit back to create a glare.

The Mueller matrices for left and right circular polarisers are defined below:

$$\begin{aligned}
 M_{left} &= \frac{1}{2} \begin{bmatrix} 1 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 1 \end{bmatrix} \\
 M_{right} &= \frac{1}{2} \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix}
 \end{aligned} \tag{5.7}$$

5.1.3 Reflectors

We mentioned previously that unpolarised light becomes partially linearly polarised upon reflection off a surface, due to Fresnel reflection effects. Such effects are represented, in Mueller calculus, by the combined effects of a linear diattenuator and linear retarder of phase δ :

$$M_r = \frac{1}{2} \begin{bmatrix} \frac{R_{\perp} + R_{\parallel}}{2} & \frac{R_{\perp} - R_{\parallel}}{2} & 0 & 0 \\ \frac{R_{\perp} - R_{\parallel}}{2} & \frac{R_{\perp} + R_{\parallel}}{2} & 0 & 0 \\ 0 & 0 & \sqrt{R_{\parallel}R_{\perp}} \cos \delta & \sqrt{R_{\parallel}R_{\perp}} \sin \delta \\ 0 & 0 & -\sqrt{R_{\parallel}R_{\perp}} \sin \delta & \sqrt{R_{\parallel}R_{\perp}} \cos \delta \end{bmatrix} \tag{5.8}$$

where R_{\parallel} and R_{\perp} are the parallel (resp. perpendicular) reflectance coefficients as predicted by Fresnel equations (eq. (2.4)) and δ is the relative phase between the parallel and perpendicular polarised components. For dielectrics, δ is a simple step-edge function of the angle of incidence: $\delta = \pi$ for any angle of incidence before Brewster angle (θ_B) and 0 otherwise. For conductors, it is a complex function of the angle of incidence [87]. Here, the input and output coordinate systems are defined with respect to the plane of incidence containing the ray of light incident at $\vec{\omega}_i$, surface normal \vec{n} and direction of perfect reflection $\vec{\omega}_r$.




Figure 5.4: **Reflector:** Geometry of a general reflecting optical element rotated by an angle $\phi_{\vec{n}}$ from its local coordinate system.

As with linear polarisers, reflectors can be rotated with respect to their local coordinate system:

$$M'_r(\phi) = M_{rot}(\phi)M_rM_{rot}(-\phi) \quad (5.9)$$

where ϕ is the inclination angle of the plane of incidence with respect to the local up vector (fig. 5.4).

5.2 Polarisation imaging

The goal of polarisation imaging is to recover the optical properties of a surface's material from observations of its reflected Stokes parameters. In particular, polarised light from reflection provides useful cues for shape estimation. A common technique to measure the complete set of parameters of the reflected Stokes (\vec{s}_r) is to take three observations of the sample under study with a linear polariser rotated at three different orientations in front of a camera and an additional measurement with a circular polariser (assumed right-handed). By direct applications of the theory presented in section 5.1, we explain why 0° , 45° and 90° are commonly employed as the orientations of the linear polariser when measuring linear reflected

Stokes parameters [62, 71, 72]:

$$\begin{aligned}
 \vec{s}_o(0^\circ) &= \left[\frac{s_{r,0} + s_{r,1}}{2}, \frac{s_{r,0} + s_{r,1}}{2}, 0, 0 \right]^T \\
 \vec{s}_o(45^\circ) &= \left[\frac{s_{r,0} + s_{r,2}}{2}, 0, \frac{s_{r,0} + s_{r,2}}{2}, 0 \right]^T \\
 \vec{s}_o(90^\circ) &= \left[\frac{s_{r,0} - s_{r,1}}{2}, \frac{s_{r,1} + s_{r,0}}{2}, 0, 0 \right]^T \\
 s_{o,righto} &= \left[\frac{s_{r,0} + s_{r,3}}{2}, 0, 0, \frac{s_{r,0} + s_{r,3}}{2} \right]^T
 \end{aligned} \tag{5.10}$$

where $s_o(\vec{\phi}_o)$ indicates the output Stokes vector as observed through a linear polariser rotated at angle ϕ_o with the camera's coordinate system as the frame of reference. The superscript T indicates matrix transposition. From eq. (5.10), it is trivially shown that these four measurements are well suited to recover the complete set of reflected Stokes parameters as:

$$\begin{aligned}
 s_{r,0} &= s_{o,0}(0^\circ) + s_{o,0}(90^\circ) \\
 s_{r,1} &= s_{o,0}(0^\circ) - s_{o,0}(90^\circ) \\
 s_{r,2} &= 2s_{o,0}(45^\circ) - s_{r,0} \\
 s_{r,3} &= 2s_{o,righto} - s_{r,0}
 \end{aligned} \tag{5.11}$$

It is worth noting that the first parameter of the reflected Stokes vector ($s_{r,0}$) can be obtained from any combination of measurements taken at orthogonal directions of the linear polariser. If we denote ϕ_o an arbitrary rotation of the

polariser, we obtain:

$$\begin{aligned}
s_{o,0}(\phi_o) + s_{o,0}(\phi_o + \frac{\pi}{2}) &= \frac{s_{r,0} + s_{r,1} \cos(2\phi_o) + s_{r,2} \sin(2\phi_o)}{2} \\
&+ \frac{s_{r,0} + s_{r,1} \cos(2(\phi_o + \frac{\pi}{2})) + s_{r,2} \sin(2(\phi_o + \frac{\pi}{2}))}{2} \\
&= s_{r,0} + \frac{s_{r,1} \cos(2\phi_o) + s_{r,2} \sin(2\phi_o)}{2} \\
&+ \frac{-s_{r,1} \cos(2\phi_o) - s_{r,2} \sin(2\phi_o)}{2} \\
&= s_{r,0}
\end{aligned} \tag{5.12}$$

The reflected Stokes vector \vec{s}_r takes a different form depending on the polarisation state of the incident illumination. Most literature on polarisation (section 2.3) has focused on resolving SFP under unpolarised incident illumination, by measuring the first three components of the reflected Stokes vector as per the protocol described above. In the following, we derive the theoretical foundations for polarisation imaging under the previously unstudied case of partially linearly polarised incident illumination (section 5.3.1) and show it to be a generalisation of the well-studied problem of polarisation imaging under the unpolarised world assumption. These derivations constitute the theoretical basis on which we build our polarisation imaging approach for reflectometry under natural illumination (chapter 6), known to be partially linearly polarised (in general) due to single-scattering of light by molecules in the atmosphere [83, 84].

5.3 Image formation model

5.3.1 Partially linearly polarised incident illumination




Figure 5.5: **Polarisation imaging geometry:** The angle of polarisation ψ_i of the incident beam of light is relative to the local coordinate system spanned by $[\vec{E}_\perp, \vec{E}_\parallel, \vec{\omega}_i]$. The outgoing coordinate system is defined by $[\vec{x}, \vec{y}, \vec{\omega}_o]$, the local coordinate system of the camera. ϕ_o is the angle of rotation of the linear polariser in front of the camera and ϕ is the angle between the camera's \vec{x} -axis and the direction perpendicular to the plane of incidence (\vec{E}_\perp).

In Stokes formalism, partially linearly polarised light is obtained from eq. (5.1) when the fourth component of the Stokes vector is 0. This corresponds to $\chi = \frac{\pi}{2}, \forall k \in \mathbb{Z}$. We thus define our incident Stokes vector for partially linearly polarised illumination as:

$$\vec{s}_i(\vec{\omega}_i) = \begin{bmatrix} L_i(\vec{\omega}_i) \\ \pm L_i(\vec{\omega}_i) \mathcal{P}_i(\vec{\omega}_i) \cos(2\psi_i(\vec{\omega}_i)) \\ \pm L_i(\vec{\omega}_i) \mathcal{P}_i(\vec{\omega}_i) \sin(2\psi_i(\vec{\omega}_i)) \\ 0 \end{bmatrix} \quad (5.13)$$

where the sign of the second and third parameters is the same and depends on the sign of $\cos(2\chi)$. The angle of polarisation (ψ_i) is defined with respect to the right-handed local coordinate system spanned by $\vec{\omega}_i$ the direction of incidence,

\vec{E}_\perp the direction orthogonal to the plane of incidence and \vec{E}_\parallel the direction parallel to the plane of incidence (fig. 5.5). Note that for sky illumination, all parameters of eq. (5.13) (L_i , \mathcal{P}_i and ψ_i) depend on the incident direction ($\vec{\omega}_i$) in a complex pattern that depends on the relative position of the sun in the sky, as predicted by Rayleigh sky model. For brevity however, we will omit the dependence on $\vec{\omega}_i$ in the remainder of this thesis.

Upon reflection off a surface, the polarisation state of an incident beam of light is changed differently by diffuse and specular reflection. It is commonly assumed that diffuse reflection depolarises light and we adopt this assumption in our work. Therefore, by application of eq. (5.4), the intensity observed at any orientation of the linear polariser for diffuse reflection (assuming Lambertian reflection, eq. (2.25)) is:

$$I_d(\phi_o) = \frac{1}{2} \frac{\rho_d}{\pi} L_i \quad (5.14)$$

where ρ_d is the material's diffuse albedo.

On the other hand, specular reflection changes the polarisation state of incident light according to eq. (5.9):

$$\begin{aligned} \vec{s}_r &= M_{rot}(\phi) M_r \vec{s}_i \\ &= L_i \left[\begin{array}{c} \frac{R_\perp + R_\parallel}{2} + \mathcal{P}_i \frac{R_\perp - R_\parallel}{2} \cos 2\psi_i \\ \frac{R_\perp - R_\parallel}{2} \cos 2\phi + \mathcal{P}_i \left(\frac{R_\perp + R_\parallel}{2} \cos 2\phi \cos 2\psi_i - \sqrt{R_\perp R_\parallel} \sin 2\phi \sin 2\psi_i \cos \delta \right) \\ \frac{R_\perp - R_\parallel}{2} \sin 2\phi + \mathcal{P}_i \left(\frac{R_\perp + R_\parallel}{2} \sin 2\phi \cos 2\psi_i + \sqrt{R_\perp R_\parallel} \cos 2\phi \sin 2\psi_i \cos \delta \right) \\ -\mathcal{P}_i \sqrt{R_\perp R_\parallel} \sin 2\psi_i \sin \delta \end{array} \right] \quad (5.15) \end{aligned}$$

Note that we have dropped the right-most rotation matrix ($M_{rot}(-\phi)$) here, because the incident Stokes vector is directly expressed in the rotated local coordinate system spanned by $[\vec{E}_\perp, \vec{E}_\parallel, \vec{\omega}_i]$ (fig. 5.5).

Finally, the intensity profile of \vec{s}_r observed through a rotated linear polariser

is obtained by the inner product of the first row of eq. (5.4) with eq. (5.15):

$$\begin{aligned}
I_s(\phi_o) = & L_i \left(\frac{R_{\perp} + R_{\parallel}}{4} + \frac{R_{\perp} - R_{\parallel}}{4} \cos(2(\phi_o - \phi)) \right) \\
& + L_i \left(\frac{R_{\perp} - R_{\parallel}}{4} + \frac{R_{\perp} + R_{\parallel}}{4} \cos(2(\phi_o - \phi)) \right) \mathcal{P}_i \cos(2\psi_i) \\
& + L_i \frac{\sqrt{R_{\perp} R_{\parallel}}}{2} \sin(2(\phi_o - \phi)) \mathcal{P}_i \sin(2\psi_i)
\end{aligned} \tag{5.16}$$

Our complete image formation model is obtained by addition of eq. (5.14) and eq. (5.16):

$$\begin{aligned}
I(\phi_o) = & I_d(\phi_o) + I_s(\phi_o) \\
= & \frac{\rho_d L_i}{2\pi} + L_i \left(\frac{R_{\perp} + R_{\parallel}}{4} + \frac{R_{\perp} - R_{\parallel}}{4} \cos(2(\phi_o - \phi)) \right) \\
& + L_i \left(\frac{R_{\perp} - R_{\parallel}}{4} + \frac{R_{\perp} + R_{\parallel}}{4} \cos(2(\phi_o - \phi)) \right) \mathcal{P}_i \cos(2\psi_i) \\
& + L_i \frac{\mathcal{P}_i \sin(2\psi) \sin(2(\phi_o - \phi)) \sqrt{R_{\perp} R_{\parallel}}}{2}
\end{aligned} \tag{5.17}$$

5.3.2 Unpolarised incident illumination

By following a similar approach to section 5.3.1, it is trivial to obtain the counterparts to Equations (5.14) to (5.17) under unpolarised incident illumination where the input Stokes vector is instead $\vec{s}_i = [L_i, 0, 0, 0]^T$. Equations (5.14) and (5.17) remain identical while \mathcal{P}_i is now 0 in eqs. (5.13) and (5.16), such that the expression

for polarisation from specular reflection simplifies to:

$$\vec{s}_r = L_i \begin{bmatrix} \frac{R_\perp + R_\parallel}{2} \\ \frac{R_\perp - R_\parallel}{2} \cos(2(\phi)) \\ \frac{R_\perp - R_\parallel}{2} \sin(2(\phi)) \\ 0 \end{bmatrix} \quad (5.18)$$

and the intensity through a rotated linear polariser (referred to as the Transmitted Radiance Sinusoid (TRS)) simplifies to (with diffuse reflection term):

$$I_s(\phi_o) = \frac{I_{max} + I_{min}}{2} + \frac{I_{max} - I_{min}}{2} \cos(2(\phi_o - \phi))$$

where (5.19)

$$I_{max} = \frac{\rho_d L_i}{2\pi} + \frac{L_i R_\perp}{2} \text{ and } I_{min} = \frac{\rho_d L_i}{2\pi} + \frac{L_i R_\parallel}{2}$$

5.3.3 Discussion

Upon comparison of the expressions obtained under partially polarised incident illumination (eqs. (5.13) and (5.16)) and their counterparts under unpolarised incident illumination (eqs. (5.18) and (5.19)), it is easy to understand why most prior work on Shape from Polarisation (SfP) [67, 88, 71, 72] has been carried under the unpolarised world assumption. In these conditions, the angle ϕ is easily obtained either from eq. (5.18) as [71]:

$$\phi = \frac{1}{2} \arctan \frac{s_{r,2}}{s_{r,1}} \quad (5.20)$$

or by fitting observations through a linear polariser oriented a three or more positions to eq. (5.19) [67, 89, 88, 72]. This in turn provides cues for Shape from Polarisation (SfP).

Under partially polarised incident illumination however, it is not clear in

general how the angle ϕ can be easily obtained, as the reflected Stokes vector now also depends on the polarisation state of the incident illumination, to the extent that some researchers [64] have stated incident linear polarisation as a fundamental limitation of SFP.

On close inspection of eq. (5.13) however, we argue that those limitations can be relaxed under certain assumptions. First, for dielectric materials, the component of reflection parallel to the plane of incidence gets completely transmitted (i.e. $R_{\parallel} = 0$) at a particular angle of incidence called Brewster angle (θ_B). Under this assumption, eq. (5.17) simplifies to:

$$I_{\theta_B}(\phi_o) = \frac{\rho_d L_i}{2\pi} + \frac{L_i R_{\perp}}{4} \cos(2(\phi_o - \phi)) + \frac{L_i R_{\perp} \mathcal{P}_i \cos(2\psi_i)}{4} \cos(2(\phi_o - \phi)) \quad (5.21)$$

which can be rearranged in the common form of the TRS (eq. (5.19) (plus diffuse term)):

$$I_{\theta_B}(\phi_o) = \frac{I_{max} + I_{min}}{2} + \frac{I_{max} - I_{min}}{2} \cos(2(\phi_o - \phi))$$

where

$$I_{max} = \frac{\rho_d L_i}{2\pi} + L_i \frac{(1 + \mathcal{P}_i \cos(2\psi_i)) R_{\perp}}{2} \quad (5.22)$$

$$I_{min} = \frac{\rho_d L_i}{2\pi}$$

We further observe that eq. (5.17) simplifies greatly when $\sin(2\psi_i) = 0$ which corresponds to the incident illumination being either horizontal (i.e. $\psi_i = 0$) or vertical (i.e. $\psi_i = \frac{\pi}{2}$). In these cases, eq. (5.17) can be written in a similar way as

eq. (5.19) (with diffuse reflection term):

$$I(\phi_o) = \frac{I_{\perp} + I_{\parallel}}{2} + \frac{I_{\perp} - I_{\parallel}}{2} \cos(2(\phi_o - \phi))$$

where

$$\begin{aligned} I_{\perp} &= \frac{\rho_d L_i}{2\pi} + L_i \frac{(1 + \mathcal{P}_i \cos(2\psi_i)) R_{\perp}}{2} \\ I_{\parallel} &= \frac{\rho_d L_i}{2\pi} + L_i \frac{(1 - \mathcal{P}_i \cos(2\psi_i)) R_{\parallel}}{2} \end{aligned} \quad (5.23)$$

Note that here, we define the TRS in terms of I_{\perp} and I_{\parallel} instead of I_{max} and I_{min} since the way I_{\perp} and I_{\parallel} compare to each other depends on the sign of $\cos(2\psi_i)$.

The expressions in eqs. (5.22) and (5.23) means that, in principle, one should recover the incident polarisation in order to recover R_{\perp} and R_{\parallel} . However, in section 6.2.3, we will discuss the challenges in recovering the incident polarisation, as well as practical guidelines to overcome those challenges. Also note that eqs. (5.22) and (5.23) are generalisations of eq. (5.19) where the latter is obtained by setting $\mathcal{P}_i = 0$.

Finally, the derivations in this chapter were made under the assumption of perfect mirror reflection and point light illumination. To account for reflection from rough surfaces under spherical illumination (as is the case under natural illumination), we model specular reflection according to the microfacet theory (eq. (2.17)). Equations (5.14) and (5.16) thus become:

$$\begin{aligned} I'_d(\phi_o) &= \int_{\Omega^+} I_d(\phi_o)(\vec{n} \cdot \vec{\omega}_i) d\vec{\omega}_i \\ I'_s(\phi_o) &= \int_{\Omega^+} I_s(\phi_o) f_s(\vec{\omega}_i, \vec{\omega}_o)(\vec{n} \cdot \vec{\omega}_i) d\vec{\omega}_i \end{aligned} \quad (5.24)$$

where $f_s(\vec{\omega}_i, \vec{\omega}_o)$ is a Cook-Torrance microfacet BRDF with a GGX (eq. (2.22)) distribution term, forming a narrow lobe around the reflection vector. Within the extent of the specular lobe, the incident polarisation can be assumed constant, as the polarisation field typically varies smoothly over the sky [90]. The derivations in

this chapter inform us on how to design a measurement protocol for reflectometry in general outdoors conditions (chapter 6).

Chapter 6

Polarisation Imaging

Reflectometry In-the-wild

Based on our theoretical derivations from chapter 5 for polarisation imaging under partially linearly polarised incident illumination, we derive a measurement protocol for reflectometry “in-the-wild” to recover high-resolution, spatially-varying reflectance maps for planar, isotropic dielectric surfaces based on commodity hardware.

In particular, we propose to image a sample at a minimum of three vantage points: one observation should be made close to normal incidence in order to have a canonical frame of reference for data registration, while two more observations should be taken close to Brewster angle of incidence, roughly orthogonal to each other, in order for eq. (5.21) to apply. Our measurement protocol is further supported by the recent work of Nielsen et al. [91] who show measurements near Brewster angle (around 60° for dielectrics) to be nearly optimal for reflectance estimation restricted to a single measurement.

The main reason for limiting the number of measurements to only 3 view points stems from a practical point of view: some of our samples (figs. 6.7 and 6.8, first row) were captured on busy walking paths where the capture process was interrupted multiple times by passers by.

The rest of this chapter is organised as follows. We first describe our main

measurement setup, employing an entry level DSLR camera (section 6.1.1), which we designed to be lightweight and accessible. We then give a quick overview of our calibration step (section 6.1.2), before providing a step-by-step analysis of our pre-processing pipeline for polarisation imaging (section 6.1.3). We continue our discussion with our reflectance recovery pipeline (section 6.2), for which we present results in section 6.3. We further give a thorough analysis of those results which we show to be on-par with other previous methods in a more controlled environment. Finally in section 6.4, we present an even lighter weight hand-held acquisition setup using a mobile device as the primary imaging system.

6.1 Digital Single-Lens Reflex (DSLR) setup

6.1.1 Data acquisition




Figure 6.1: **Principal polarisation imaging setup:** We employ commodity photography equipment often used for image-based lighting applications [4].

Our main measurement setup is designed around commodity photography equipment, often used for image-based lighting applications [4]:


- **A DSLR camera:** In our experiments, we used a 17.9 MP Canon EOS 650D camera with an 18-55 mm lens to which we mounted a 58 mm rotatable linear polariser from Edmund Optics. Throughout the capture process, the camera is mounted on a heavy duty tripod for stability.

- **A calibration target:** We used an X-Rite ColorChecker® chart placed next to the sample for radiometric calibration.
- **A chrome ball:** To record the incident illumination, we used a chrome ball generally used as garden decoration.

Figure 6.1 shows a typical arrangement of our DSLR-based capture setup for outdoor measurements, where we manually rotate the linear polariser in front of the camera at three marked orientations (0° , 45° , 90° , marked on the polariser) to image the polarisation of light reflected off the sample, as per eq. (5.10).

Our measurement protocol (fig. 6.2) proceeds as follows: we first image the s_0 component of the reflected Stokes field of the sample close to normal incidence, to provide a canonical frame of reference for data registration. While this could have been done by photographing the sample without polariser, we found it less cumbersome to leave the polariser on at all time during capture and simply image the 0° and 90° orientations only at normal incidence. We then proceed to the complete measurement of the linear Stokes parameters for two roughly orthogonal views close to Brewster angle of incidence. While in principle, eqs. (5.21) and (5.22) are only true at the exact Brewster angle of incidence, we found that being in a $\pm 15^\circ$ window around Brewster angle is sufficient, in practice, for good qualitative measurements. This will be discussed in more details in section 6.3.2. Typical measurements took us around 5 minutes per sample when uninterrupted and up to 20 minutes when interrupted by passers by, as was the case for our “drain cover” sample (figs. 6.7 and 6.8, first row).

We propose a simple practical guideline for finding the best near-Brewster measurement in practice as follows: with the polariser oriented at 90° (i.e vertical) with respect to the camera’s \vec{x} -axis pointed at the sample, adjust the camera’s height until minimum transmission through the polariser is achieved. In this configuration, the camera is seeing the sample at a near-Brewster angle defined from the sample’s mean up-vector. The rationale behind this idea is that given the planar nature of our samples, the plane of incidence is mainly vertical and the horizontally reflected light at Brewster angle is therefore cross-polarised with the



Normal incidence Near-Brewster incidence

Figure 6.2: **Polarisation imaging reflectometry in the wild:** Measurement protocol. We acquire a set of HDR sequences near normal incidence and close to Brewster angle of incidence.

vertical polariser on the camera.

Each near-Brewster measurement consists of multiple frames per orientation of the polariser, where each frame is taken at a different exposure level in order to recover HDR maps of the reflected Stokes parameters. To mitigate the effects of changes in illumination at each oblique views, we used the Auto-Exposure Bracketting (AEB) setting of our camera for rapid acquisition of the HDR sequence.

For each view, the HDR sequences captured at 0° , 45° and 90° are then combined into the linear sRGB HDR radiance maps $I(0^\circ)$, $I(45^\circ)$ and $I(90^\circ)$ respectively, using `pfstools` [92]. While many other HDR packages exist that could replace `pfstools`, we chose the latter for its command line interface which allows easy batch processing through custom scripts.



Figure 6.3: **Semi-automatic registration tool:** The user selects four matching corners on each frame and our tool automatically generates a regular grid of matched points ready to pass as input to VisualSfM.

6.1.2 Calibration

Radiometric calibration

As for any data driven approach for reflectometry, calibration of the acquired data is essential before any further processing. Similar to our mobile reflectometry approaches (part II), we rely on the presence of the X-Rite ColorChecker® next to the sample for radiometric calibration of both the sample’s response and light probe captured by the chrome ball. After assembling the HDR frames ($I(0^\circ)$, $I(45^\circ)$ and $I(90^\circ)$), we scale their overall brightness such that the white point of the calibration target reads an intensity of $[0.45, 0.45, 0.45]$. Note that the linear sRGB value for this patch is $[0.9, 0.9, 0.9]$ but eq. (5.14) predicts that the observed intensity through a linear polariser for diffusely reflected light is halved.

Geometric calibration

For surface normal estimation (section 6.2.1) and roughness estimation through inverse rendering (section 6.2.3), we further require camera pose estimation. For this purpose, we employ Wu’s VisualSfM [93], a GUI-based Structure from Motion

(SfM) software package. In some cases, we found VisualSfM to have difficulties finding enough feature correspondences for geometric calibration. In such cases, we thus developed a simple tool to help bootstrap the calibration process, where the user manually selects matching rectangles on each image in the sequence (fig. 6.3). The tool then automatically generates a regular grid of feature points which are exported as Scale-invariant Feature Transform (SIFT) descriptors [94], ready to be imported into VisualSfM’s ecosystem.

From VisualSfM, we obtain camera poses and tracks (set of correspondences between 3D world coordinates and 2D projected points on the image plane), which are refined through bundle adjustment [95]. From those tracks, we compute homography matrices to register the acquired data to the canonical view captured at normal incidence, and perform all subsequent calculations in image space of the canonical view.

6.1.3 Polarisation processing pipeline




Figure 6.4: **TRS fitting**: For each near-Brewster view, we compute a per-pixel fit of eq. (5.22) to the acquired data.

Given a set of HDR observations through a linear polariser oriented at $\phi_o = 0^\circ, 45^\circ, 90^\circ$ for two roughly orthogonal observations new Brewster angle of incidence, we start by fitting eq. (5.22) per-pixel (in the least-squares sense) to each set of observations. While, in principle, any non-linear optimisation package

could be used to perform this task, we found such approach to be inefficient in practice, with fitting times in the order of two hour at our camera's resolution, when fitting with the curve fitting routine from scipy's optimisation package [?].

Instead, we note that by rearranging the terms in eq. (5.22) as in eq. (6.1), we obtain a linear problem of the form $Ax = b$ which can be solved for very efficiently through Singular Value Decomposition (SVD), bringing the fitting times under ten seconds.

$$\underbrace{I_{\theta_B}(\phi_o; x, y)}_b = \underbrace{\begin{bmatrix} 1 & \cos 2\phi_o & \sin 2\phi_o \end{bmatrix}}_A \underbrace{\begin{bmatrix} \frac{I_{max} + I_{min}}{2} \\ \frac{I_{max} - I_{min}}{2} \cos 2\phi \\ \frac{I_{max} - I_{min}}{2} \sin 2\phi \end{bmatrix}}_x \quad (6.1)$$

The per-pixel parameters of the TRS, can then be obtained from the intermediate result of this linear formulation ($\hat{x} = \begin{bmatrix} x_1 & x_2 & x_3 \end{bmatrix}^T$) as:

$$\begin{aligned} \hat{I}_{max} &= x_1 + \sqrt{x_2^2 + x_3^2} \\ \hat{I}_{min} &= x_1 - \sqrt{x_2^2 + x_3^2} \\ \hat{\phi} &= \frac{1}{2} \arctan \frac{x_3}{x_2} \end{aligned} \quad (6.2)$$

Figure 6.4 presents TRS parameter maps for one of our samples, obtained from the fitting protocol described above. These maps constitute the building blocks from which we derive our reflectance estimation pipeline (section 6.2).

6.2 Reflectance recovery




Figure 6.5: **Polarisation imaging reflectometry:** Example reflectance maps recovered for a permanent on-site specular “drain cover”, captured on a busy sidewalk close to Imperial College’s campus.

6.2.1 Surface normal estimation

We formulate our normal estimation in a multi-view SfP framework, akin to that of Wolff [67] and Miyazaki et al. [88]. Consider the projection of \vec{E}_\perp on the image plane. From (fig. 5.5) it is clear that this vector, which we denote \vec{b} , is defined as:

$$\vec{b} = \begin{bmatrix} \cos \hat{\phi} & \sin \hat{\phi} & 0 \end{bmatrix}^T \quad (6.3)$$

where $\hat{\phi}$ is obtained from eq. (6.1). The vector \vec{b} is, by definition, orthogonal to the surface normal \vec{n} , therefore the following expression always holds true:

$$(\vec{b} \cdot \vec{n}) = 0 \quad (6.4)$$

At each view point, knowing \vec{b} essentially constrains the surface normal to lie in the plane of incidence. Combining at least 2 observations of the TRS phase angle $\hat{\phi}$ from different view points therefore provides enough constraints to fully recover \vec{n} .

Given our two observations of the TRS phase angle ($\hat{\phi}_0$ and $\hat{\phi}_1$ at two view-points close to Brewster angle of incidence (fig. 6.4c), with their respective camera

rotation matrices in world coordinates (R_0 and R_1) obtained from VisualSfM, the surface normal is recovered by solving the following linear problem:

$$\underbrace{\begin{bmatrix} R_0^T \vec{b}_0 \\ R_1^T \vec{b}_1 \end{bmatrix}}_{2 \times 3} \begin{bmatrix} n_x \\ n_y \\ n_z \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad (6.5)$$

Similarly to eq. (6.1), the above problem can be solved for efficiently using SVD. Figure 6.5c shows the normal map we recovered for our “drain cover” dataset.

6.2.2 Diffuse albedo estimation

In our framework, diffuse albedo estimation is straightforward from the observation of \hat{I}_{min} . From eqs. (5.22) and (5.24), we have:

$$\begin{aligned} \hat{I}_{min} &= \frac{\rho_d}{2\pi} \underbrace{\int_{\Omega^+} (\vec{n} \cdot \vec{\omega}_i) L_i d\vec{\omega}_i}_{\pi} \\ &= \frac{\rho_d}{2} \end{aligned} \quad (6.6)$$

where the integral part is simplified because of our radiometric calibration step (section 6.1.2). The diffuse albedo is thus simply obtained as $\rho_d = 2\hat{I}_{min}$. In principle, any of the two views close to Brewster incidence can be used to estimate the diffuse albedo. However in practice, we found that despite our best efforts to be as close as possible to Brewster angle as per the protocol described in section 6.1.1, one of the two views may, in some cases, show slightly better specular cancellation (see fig. 6.4b for an example). We thus compute $2\hat{I}_{min}$ for each view and set the diffuse albedo (see fig. 6.5a) as the minimum of the two.

6.2.3 Specular component estimation

Influence of incident polarisation




Figure 6.6: Change in reflected radiance $s_{r,0}$ due to incident polarised illumination ($DOP = 80\%$) is very similar for stainless steel (solid plots) and a dielectric ($\eta = 1.5$, dotted plots) around Brewster angle of incidence. The colours indicate three different angles of polarisation w.r.t. the plane of incidence (Red: $\psi_i = 0^\circ$, Green: $\psi_i = 45^\circ$, Blue: $\psi_i = 90^\circ$).

As pointed out at the end of section 5.2, the measured reflected intensities ($s_{r,0}$) depend on the state of polarisation of the incident illumination which, in principle, should be recovered in order to solve for eqs. (6.7) and (6.9) described thereafter for specular reflectance and roughness estimation. However, in practice, recovering the incident Stokes field from a mirror ball is challenging as reflection from metals becomes elliptical, thus requiring measurements with both a linear and circular polariser.

In our experiments, we found that constantly swapping a linear and circular polariser in front of the camera was cumbersome for outdoors measurements and defeated the purpose of simplicity of our method. Inspired by the work Legendre et al. [96], we thought about employing a black shiny dielectric sphere instead, in order to solve for the issue of reflected elliptical illumination. This did not work either however, as recovering the incident Stokes field requires inverting the 4x4 Mueller matrix of the dielectric sphere, which becomes singular at Brewster angle due to total transmission of the p-polarised component.

Instead, we propose a practical solution based on the observation that the incident radiance recorded using the mirror ball already encodes the modulation

of intensity in the $s_{r,0}$ component due to the incident partial linear polarisation (fig. 6.6). This is why light probes captured outdoors exhibit darker and brighter sections in the sky due to polarisation effects. Figure 6.6 shows that this change in intensity of reflected light observed on a stainless steel mirror ball is very similar to that on a dielectric around Brewster angle. With this in mind, we propose first order approximations to solve for specular reflectance and roughness.

Specular reflectance

From the same view as that used for diffuse albedo estimation, we obtain a specular only observation by subtracting half the diffuse albedo to \hat{I}_{max} :

$$\begin{aligned}\hat{I}_{s,max} &= \hat{I}_{max} - \frac{\rho_d}{2} \\ &= \int_{\Omega^+} \frac{R_{\perp}(\mathcal{P}_i \cos 2\psi_i + 1)}{2} f_s(\vec{\omega}_i, \vec{\omega}_o) L_i(\vec{n} \cdot \vec{\omega}_i) d\vec{\omega}_i\end{aligned}\tag{6.7}$$

This diffuse-free image encodes R_{\perp} up to a scale factor, and knowledge of R_{\perp} is sufficient to recover the material's index of refraction as demonstrated by Ghosh et al. [62], and subsequently its reflectance at normal incidence ($F(0^\circ)$, see fig. 6.5b) for use with Schlick's approximation [19] to model Fresnel reflectance:

$$\begin{aligned}\eta &= \sqrt{\frac{1 + \sqrt{R_{\perp}}}{1 - \sqrt{R_{\perp}}}} \\ F(0^\circ) &= \frac{(\eta - 1)^2}{(\eta + 1)^2}\end{aligned}\tag{6.8}$$

In practice, recovering R_{\perp} is challenging as it depends both on the surface's specular roughness and polarisation state of incident illumination, which is not easily obtained, even when capturing a light probe for the reasons stated earlier.

Therefore, we instead propose a practical, template-based approximation where we recover the scale factor from the observation of the diffuse-free signal on the plastic casing surrounding the calibration target. The latter is made of plastic for which we borrow the index of refraction $\eta_{chart} = 1.46$ from online

sources [97, 98]. From this, we first compute the chart's perpendicular reflection coefficient at Brewster angle under uniform spherical illumination, $R_{\perp,chart}$. We then compute the scale factor between the charts measured diffuse-free intensity and $R_{\perp,chart}$ and apply the same factor to the samples diffuse-free intensity image, to obtain an estimate of R_{\perp} . The idea is that, since our samples are planar, they are roughly subject to the same incident illumination as the calibration target placed flat close to the sample.

Specular roughness

Finally, we formulate our roughness estimation stage as an inverse rendering problem where we find the value σ that best explains our observations. To do so, we solve the following least squares problem:

$$\min_{\sigma} \frac{1}{2} \sum_i \|I_{r,i} - \hat{I}_{r,i}(\sigma)\|_2 \quad (6.9)$$

where $I_{r,i}$ is our measured reflected intensity at each near-Brewster views, and $\hat{I}_{r,i}$ is rendered at each stage of the optimisation, given estimates of the diffuse and specular reflection components and normals from the previous steps, and camera poses with their respective light probes. The latter encodes the modulation of intensity in I_r due to incident partial linear polarisation. We found that this first order approximation of incident polarisation with view dependent light probes gave satisfying results in practice (see figs. 6.7 and 6.8). The values of $I_{r,i}$ are never measured directly, but recalling eq. (5.12), it was shown that the sum of observations from any two orthogonal orientations of the polariser yields $s_{r,0}$, so in particular:

$$I_r = s_{r,0} = \hat{I}_{min} + \hat{I}_{max} \quad (6.10)$$

The rendered intensities ($\hat{I}_r(\sigma)$) are computed as follows:

$$\hat{I}_r(\sigma) = \int_{\Omega^+} \left(\frac{\rho_d}{\pi} + \frac{F(\theta)G(\vec{\omega}_i, \vec{\omega}_o, \vec{n})D(\vec{\omega}_i, \vec{\omega}_o, \vec{n}, \sigma)}{4|\vec{n} \cdot \vec{\omega}_i||\vec{n} \cdot \vec{\omega}_o|} \right) L_i(\vec{n} \cdot \vec{\omega}_i) d\vec{\omega}_i \quad (6.11)$$

where L_i is obtained from each view’s light probe, ρ_d is the diffuse albedo estimated as in section 6.2.2, \vec{n} is the surface normal estimated as described in section 6.2.1, $F(\theta)$ is computed by Schlick’s approximation given our estimate of $F(0^\circ)$, and $D(\vec{\omega}_i, \vec{\omega}_o, \vec{n}, \sigma)$ is a GGX distribution term (eq. (2.22)).

6.3 Results and analysis

6.3.1 Results

To assess the validity of our method, we conducted a series of outdoor measurements of planar surfaces under varying illumination conditions. Note that most of our samples are permanent on-site structures that cannot be brought in a laboratory for measurement. Figure 6.7 shows reflectance maps (a-d) that we estimated for various datasets following the steps outlined in section 6.2, each measured under different illumination conditions (e). Photo-rendering comparisons are provided in fig. 6.8 (a,b), as well as renderings in novel light environments (c)

We chose these exemplar datasets as they exhibit very different reflectance properties. The “drain cover” (figs. 6.7 and 6.8, first row) was captured on a sidewalk near Queen’s Gate in South Kensington, which is a busy area due to the close proximity of the Natural History Museum, Victoria & Albert Museum and Hyde Park. As can be seen on the light probe (fig. 6.7, first column, e), it was captured in an environment with trees and buildings all around, showing robustness to clutter in the incident illumination.

Furthermore, it is an interesting sample as it is not strictly speaking a dielectric. Indeed, drain covers are generally made of cast iron, a composite of iron (metal) and carbon (dielectric). Nonetheless, we obtain good qualitative results for this dataset, which, from our observations, we attribute to two main factors:

1. Dielectric behaviour tends to dominate in metal-dielectric composites.
2. Outdoor metallic surfaces are subject to weathering effects causing oxidation

which adds to the dielectric-like behaviour.

Our first observation is further validated by our “red book” sample (figs. 6.7 and 6.8, second row) which is covered with a thin layer of metal-dielectric composite paint. It was captured in an open environment with partial cloud coverage. It presents a complex texture pattern which is faithfully captured by our reflectance maps, which we later (section 6.3.2) compare to two methods in controlled environments, based on commodity hardware:

1. The LCD-based reflectometry approach of Ghosh et al. combining polarisation and spherical gradient illumination [5].
2. The two-shot SVBRDF capture approach of Aittala et al. [14].

In both cases, we show good qualitative and quantitative agreement to support our method. The “red bricks” and “garden pavement” samples (figs. 6.7 and 6.8, third and fourth rows) are both diffuse dominated samples captured in very different lighting environments: the first one was captured under a fairly uniform blue sky with a mild cloud coverage while the second was captured under full overcast conditions. Note that the latter is an ideal condition in the sense that it corresponds exactly to the well-studied case of polarisation imaging under the unpolarised world assumption. It can be seen in fig. 6.7a, fourth row that the surface normals are slightly noisy in this case. This is to be expected as the sample is mostly diffuse thus giving a low Signal-to-Noise Ratio (SNR) in its specular polarisation signal. Even then, our method is still able to recover reflectance maps that go a long way towards realism (fig. 6.8, fourth row).




Figure 6.7: Reflectance maps ((a)-(d)) estimated from two views of the sample close to Brewster angle of incidence, under natural outdoors illumination (e). Our method is agnostic to the incident illumination and robust to changes in illumination during capture.

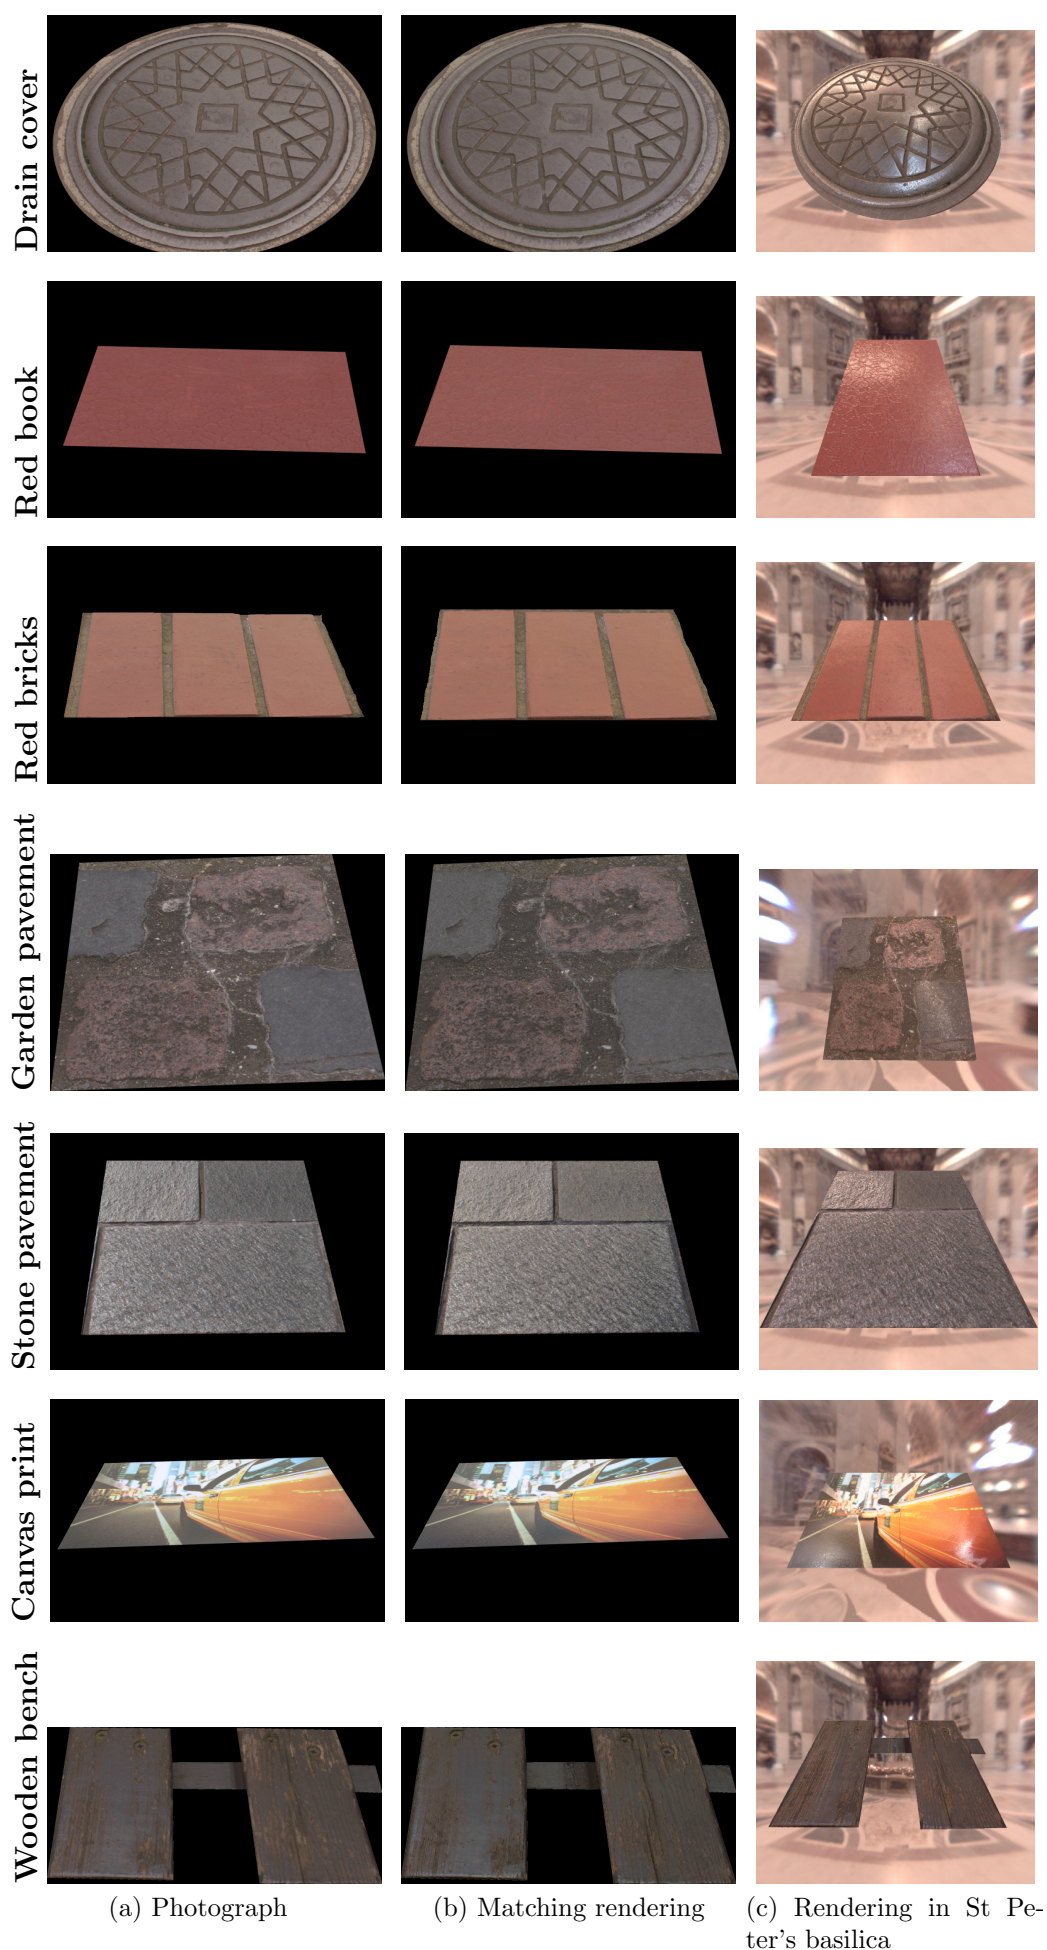


Figure 6.8: Comparisons of sample photographs (a) to matching renderings under the same incident illumination (b), as well as renderings in novel lighting environment (c).

Our “stone pavement” sample (figs. 6.7 and 6.8, fifth row) is another example of diffuse dominated material, captured on a sunny afternoon under open blue sky. The surface normals do not suffer from noise here as the signal is stronger under sun light. The principal drawback of the latter however can be seen in the specular reflectance map (fig. 6.7b, fifth row) where some directional lighting artefacts remain, due to the sun being the dominant source of illumination.

The “canvas print” (figs. 6.7 and 6.8, sixth row) shows significant specular variation over the sample, due to the varying concentrations of ink over the sample.

Finally, the “wooden bench” sample (figs. 6.7 and 6.8, last row), captured in an interior courtyard surrounded by tall buildings, shows significant variations in specularity due to varying levels of wear and tear to the layer of varnish. Some artefacts can be seen on the support panel in the middle, due to ambient occlusion which our method does not model.

Overall, our method captures reflectance maps with rich details, for samples ranging from diffuse dominated (figs. 6.7 and 6.8, third to fifth rows), to highly specular (figs. 6.7 and 6.8, first and second rows) that provide good qualitative results for photo-realistic rendering.

6.3.2 Discussion and error analysis

The main assumption for our method is that measurements be made close to Brewster angle. In the following, we provide a theoretical analysis of the Brewster measurement assumption and in particular, how close is “close to Brewster angle”. We also provide quantitative and qualitative comparisons of our reflectance maps for the “red book” sample (fig. 6.7, second row) at different times of day, to measurements carried in a controlled environment with commodity hardware, in order to assess the accuracy of our method in various lighting conditions.

Surface normal accuracy

We rely on the detection of the maximum intensity of the TRS at two different views, to provide cues to reliably estimate surface normals. We thus conducted

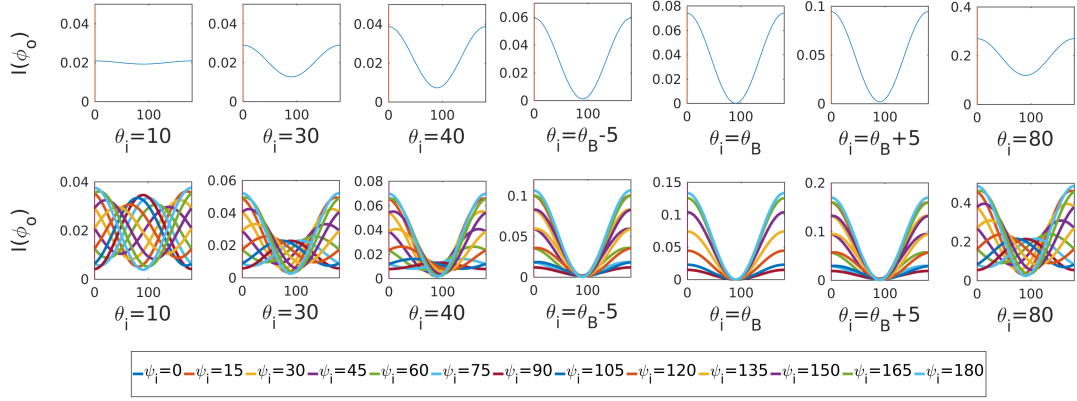


Figure 6.9: **Brewster angle measurement validation:** Simulated TRS for a glass material (Index of Refraction $\eta = 1.5$) oriented at an azimuth $\phi_{\vec{n}} = 90^\circ$. First row: Simulation under unpolarised incident illumination - the maximum of the TRS is found at $\phi_o = 0^\circ$, as expected, for any angle of incidence θ_i . Second row: Simulation under partially linearly polarised illumination with a DOP of 80%. The different colours represent different angles of polarisation ψ_i . Unlike under unpolarised incident illumination, the phase of the TRS depends on the angle of polarisation of the incident illumination. However, behaviour similar to that obtained under unpolarised illumination can be observed again at and around Brewster angle (i.e. when $\theta_i = \theta_B$).

a theoretical analysis of the changes in transmitted radiance through a linear polariser both under unpolarised illumination (fig. 6.9, first row) and partially linearly polarised illumination with a DOP of 80% which is the maximum predicted by Rayleigh sky model (fig. 6.9, second row). The latter is essentially the worst case scenario for passive outdoors reflectometry. We simulated the TRS at different angles of incidence, for a flat surface pointing at the vertical (i.e. its azimuth is at 90° from the horizon).

As can be seen in fig. 6.9 (first row), the maximum of the sinusoid is at 0° for any angle of incidence, which is exactly what we expect: the phase of the sinusoid does not depend on the angle of incidence under the unpolarised world assumption and multi-view surface normal estimation is reliable even for curved surfaces as has been demonstrated multiple times in previous work [67, 88]. Under partial linear polarisation however (fig. 6.9, second row), the phase of the sinusoid varies with the angle of polarisation of the incident illumination (coloured curves). A behaviour similar to that of the unpolarised world assumption can be observed again at Brewster angle ($\theta_i = \theta_B$). Interestingly, this still holds to some extent in a 15° window around Brewster angle where we found the estimation error on

		Ours		Ghosh et al. [5]		Aittala et al. [14]	
Surf. normals	Time Of Day	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean
	Cloudy (mid-day)	3.80°	[0, 0, 1]	5.32°	[0, 0, 1]	5.13°	[0, 0, 1]
	Sunny (10-10:30am)	8.91°					
	Sunny (3-3:30pm)	5.97°					
	Sunny (6-6:30pm)	8.97°					

Table 6.1: Statistical variation in surface normals of “red book” under different lighting conditions (left column), compared to two measurement methods employing controlled illumination.

ϕ to be less than 4° , which is acceptable for rendering applications. Furthermore, while Brewster angle can vary over large surfaces, we only found a $5 - 6^\circ$ variation over our largest sample (“canvas print”, figs. 6.7 and 6.8, fifth row), which is still within the acceptable range as per our analysis.

We further conducted a statistical analysis of surface normals estimated from our method for the “red book” (figs. 6.7 and 6.8, second row) at different times of day under an open sky, to the recent work of Ghosh et al. [5] and Aittala et al. [14] on surface reflectometry using commodity hardware. Table 6.1 summarises our analysis.

We chose the different times of day to reflect changes in polarisation of the sky light as predicted by Rayleigh sky model: at sunset and sunrise, the sky is strongly vertically polarised at the zenith in the North-South direction, while it is mostly horizontally polarised at mid-day. From the figures in table 6.1, we make the following observations: the cloudy measurements (table 6.1, first column) and sunny measurements at 3pm (table 6.1) seem to be the best conditions for data acquisition. The former makes sense as the unpolarised world assumption (ideal case) prevails. The latter comes from the fact that the measurements were carried close to the mid-day sun, which we determined to be around 1-1.30 pm on the day of the capture, where the incident illumination was mostly horizontal. From our theoretical analysis in section 5.2, this makes sense as horizontal polarisation simplifies the mathematics of the TRS to a form close to that of the unpolarised world assumption. The other two conditions at 10am (table 6.1, second row) and 6pm (table 6.1, fourth row) correspond to non ideal cases with respectively an arbitrary polarisation angle and vertical polarisation. These effects are however mitigated in part by the fact that we carry our measurements close to Brewster angle of




Figure 6.10: **Diffuse-specular separation at normal incidence:** At 6pm, the sky is strongly linearly polarised at the zenith (a), which allows for good diffuse-specular separation ((c),(d)) near normal incidence.

incidence. The 6 pm measurement is interesting for another reason however, as measurements at normal incidence provide a good way of separating diffuse and specular reflectance (fig. 6.10). This is because at that time, the band of highest degree of polarisation is directly overhead.

On the practical sides, a major challenge that arises from measurements under sunny conditions stems from the fact that the sun’s intensity largely dominates that of the sky light. This in turn creates visible directional bias, an example of which can be seen in the specular reflectance map of the “stone pavement” fig. 6.7, fifth row.

Index of refraction accuracy

For lack of measurement device capable of accurately measuring the index of refraction of the colour chart’s casing in our laboratory, we borrowed the value of η_{chart} from online sources [97, 98]. This is in line with prior work on skin reflectance [99, 58] and SfP [64, 72, 73] which have shown the index of refraction to have a very marginal impact on reflection from dielectrics. For further validation of index of refraction estimation, we also compared our estimates for the “canvas print” surface to known values of various types of inks [100], which are as follows: black 1.65 (ref. 1.6), blue & yellow 1.49 (ref. 1.53 – 1.54).

Reflectance maps comparison




Figure 6.11: **Reflectance maps comparison:** We compare the reflectance maps of the “red book” sample recovered with our method (first row) to those recovered from two methods under controlled illumination (second and third row). Our method shows good qualitative agreement with those methods.

Figure 6.11 compares the diffuse albedo (a), specular reflectance (b) and normal maps (c) obtained with the method presented in the previous sections (first row) for our “red book” sample (figs. 6.7 and 6.8, second row) to those obtained from more controlled setups: Ghosh et al. [5] (second row) and Aittala et al. [14] (third row). As can be seen, despite the challenging measurement conditions, our method produces reflectance maps that are qualitatively similar to more controlled approaches. It is interesting to note that similar to Aittala et al. [14], our method produces a rather colourless specular albedo map, which seems to disagree with the approach of Ghosh et al. [5]. In the latter case, we attribute the presence of colour in the specular reflectance map to the fact that the “red book” surface is coated with a slightly metallic paint which is not modelled with our approach which assumes pure dielectric reflectance.

We also note that compared to our method and that of Ghosh et al., the method proposed by Aittala et al. does not produce exact reflectance maps: their output well mimics the statistical properties of the material, in this example, by capturing the stochastic nature of the “red book” sample, with the exception of the diffuse albedo map (fig. 6.11a) where inter-reflections artefacts are present.

Overall, given the challenging measurement environment, our method does produce reflectance maps that go a long way towards producing compelling renderings (figs. 6.8b and 6.8c) and are on par with more controlled methods based on commodity hardware.

Finally, we also note that our method does not explicitly consider diffuse polarisation as it has been shown to have a minor contribution ($\sim 9\%$) to polarised reflectance near Brewster angle [59, 64] (fig. 2.5, red curve). Furthermore, given our day-time measurements of upwards facing planar surfaces, any resulting diffuse polarisation will be mostly p-polarised (due to in-plane transmission) and contribute to I_{min} and hence to the diffuse albedo estimate with our measurements.

Limitations

Currently, our method makes a strong assumption that the material should be isotropic, dielectric and planar for our theory (chapter 5) to apply. However, this already covers a wide variety of real world materials, and we showed in practice that some of these assumptions can be relaxed, by demonstrating reflectance acquisition for dielectric-metal composites. For materials with a strong metallic behaviour, we did encounter issues with our method, as can be seen in the zoomed in diffuse and specular maps (figs. 6.12a and 6.12b), showing the diffuse-specular separation predicted by our method for an ornate book cover. Here, the burgundy coloured pattern is a leather-like material for which diffuse-specular separation works well. On the other hand, the golden patterns are coloured with golden paint which behaves like a metal; most of the golden colour should thus be pushed to the specular reflectance map while the diffuse albedo map should have almost no colour, which

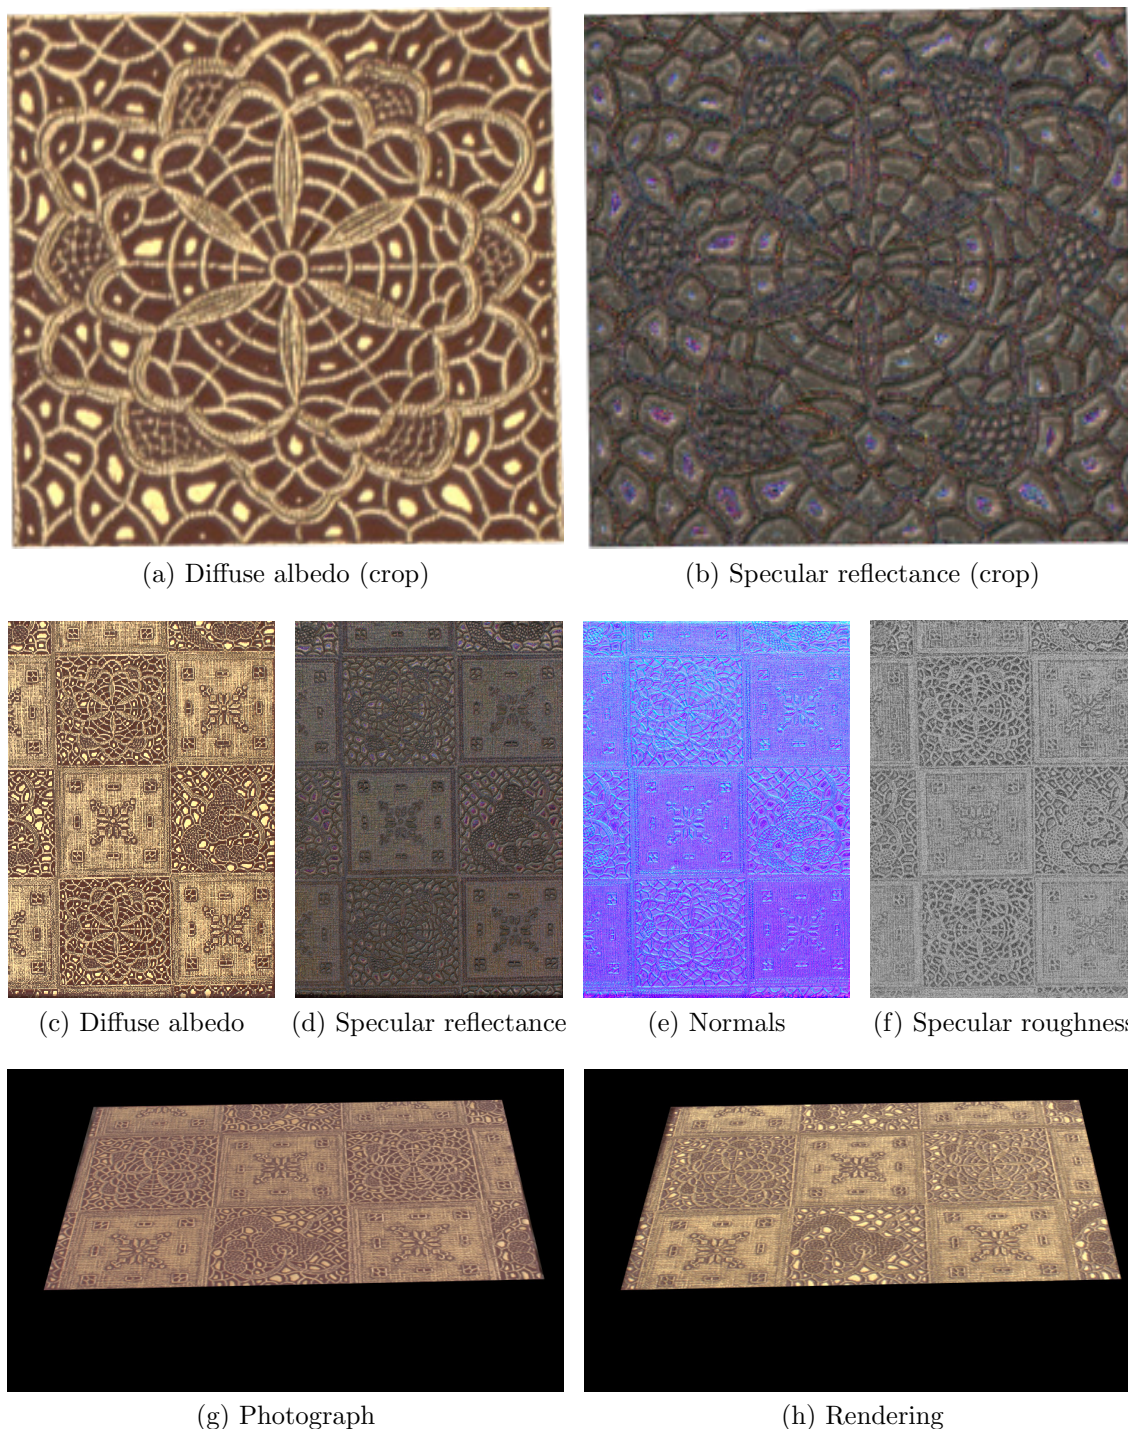


Figure 6.12: Our method has problems with accurate estimation of surface reflectance over the metallic surface of this ornate book cover.

is not the case. Furthermore, while the reflectance maps (Figures 6.12c to 6.12f) seem to present rich qualitative surface details, they suffer from high absolute error (RMS error: 28.4° angular in the normal map), which further affects the inverse rendering step, as can be seen with the rendered frame (fig. 6.12h) not matching the photograph (fig. 6.12g).

We also found our method to be sensitive to the level of occlusion of the incident illumination. Figure 6.13 (first row) shows the reflectance maps we recovered for a bas-relief mounted on an interior wall where most of the lighting comes from a small set of directions at the top right with respect to the camera (fig. 6.13a). Even in this complex arrangement, our method produces reasonable diffuse, specular and normal maps (Figures 6.13b to 6.13d), with nonetheless the following visible artefacts:

1. We first note that the different maps show some self-occlusion artefacts around regions of high curvature (nose, fingers), which do not quite agree to our planar surface assumption (figs. 6.13b to 6.13c).
2. Second, the surface normals are biased in the x-direction due to direct lighting coming only from a small localised area (fig. 6.13a).
3. Finally, we are not able to reliably estimate the specular roughness, to match photograph (fig. 6.13e) and rendering (fig. 6.13f).

Despite those limitations, we consider this example to be a part success/failure scenario which, given the complex setup, still produces reasonably believable rendered imagery (figs. 6.13f and 6.13g).

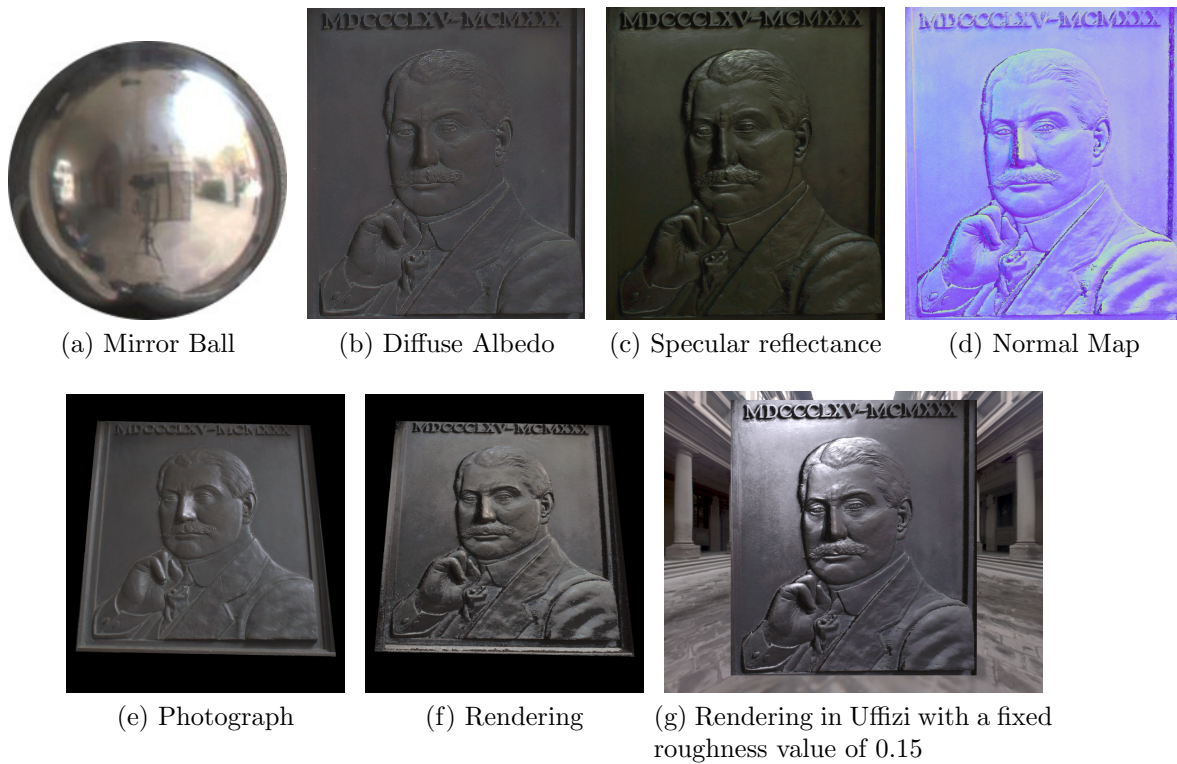


Figure 6.13: **Partial result:** Bas-relief on an interior wall of an enclosed entrance to a courtyard acquired with our method.

6.4 Mobile polarimetric setup

Besides our primary acquisition setup (section 6.1), we developed an even lighter protocol, where we replaced the DSLR camera and tripod device by a hand-held mobile device, to which we statically mounted a piece of linear polarising gel sheet. The different orientations of the polariser are then obtained by directly rotating the device while the internal sensors are queried to estimate the current orientation of the polariser. In this section, we first describe the necessary modifications to the mobile device for polarimetric measurements and the calibration process specific to the hand-held nature of the setup (section 6.4.1), followed by some results acquired with our hand-held mobile polarimetric setup (section 6.4.2).

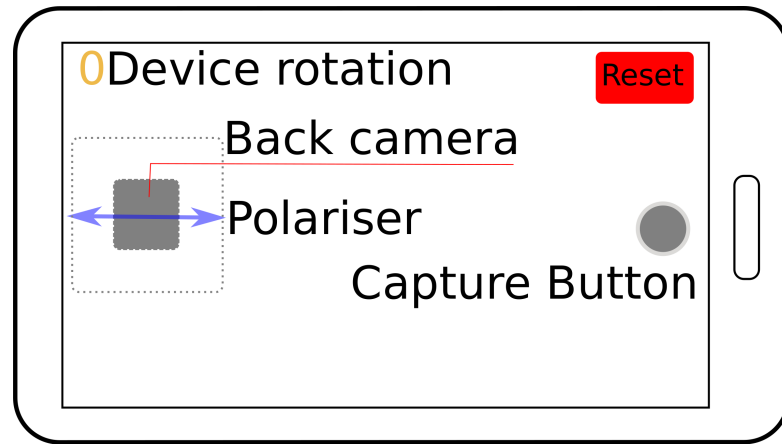


Figure 6.14: **Mobile polarimetric setup:** We statically mounted a linear polariser to a mobile phone’s back camera and developed a custom application to allow mobile polarimetric measurements. The polariser’s optical axis is oriented parallel to the device’s landscape position.

6.4.1 Acquisition pipeline

For our mobile polarimetric setup, we replaced the Canon EOS 650D DSLR camera and tripod with a Samsung Galaxy S4 mobile phone, running Android 5.0.1 “Lollipop”. Its primary (back) camera has a 13 MP resolution sensor and a 31 mm fixed focal lens with fixed $f/2.2$ aperture. While clip-on circular polarising filters are available for mobile devices nowadays, linear polarisers are currently unavailable. We thus statically mounted a laminated plastic sheet of linear polarising film from Edmund Optics in front of the lens, such that the transmission axis of the polariser is horizontal when the device is in landscape mode (fig. 6.14).

Our hand-held capture process then proceeds as follows: the user performs an in-plane rotation of the device from landscape all the way to portrait mode, while a live preview is displayed on-screen for immediate feedback. Similar to our free-form mobile surface reflectometry setup (chapter 3), we query the phone’s internal IMU sensors to estimate its relative rotation in real-time. This information is displayed in the top-left corner of the live preview to allow the user to capture the 0° , 45° and 90° orientations (fig. 6.15). Given the hand-held nature of the acquisition setup, we do not require that the user measure exactly the 0° , 45° and 90° orientations as in principle any 3 sampling points well spaced angularly are sufficient to fit the TRS. Instead, we allow the user to get as close as possible



Figure 6.15: **Mobile polarimetric acquisition:** The user rotates the device at 0° (a), 45° (b) and 90° (c) in order to sample the TRS. Notice how the intensity of the reflection on the “drain cover” changes with the orientation of the device.

to those values and save the IMU sensors’ readings together with each captured frame.

Also related to the hand-held nature of the acquisition setup, we limited the capture to a single exposure in order to minimise motion blur. We ensured that the sample and colour chart were well exposed within a single exposure, which resulted in the mirror sphere being over exposed. We were thus unable to recover the specular roughness with this setup. However, we show in section 6.4.2 that our diffuse, specular and normal maps already go a long way into creating compelling

renderings with an homogeneous, manually specified specular roughness value.

Data calibration

Unlike with a DSLR camera, with our mobile polarimetric setup, the entire scene is recorded at different orientations. Therefore, registration has to be executed not only across views but also within each viewpoint. Here, we found our semi-automatic tool for homography-based registration tool (section 6.1.2) to be helpful also for registration within each view.

The rest of the calibration process is then similar to the previous DSLR-based setup: radiometric calibration is done relative to the X-Rite ColorChecker® calibration target and camera calibration obtained from VisualSfM [93].

Processing

With the data registered, the processing pipeline is then the same as described in section 6.2. The only notable difference is that we cannot recover specular roughness as we were unable to obtain a well exposed image of the mirror ball and sample in one single exposure. Furthermore, our data calibration step assumes planar surfaces: after registration of the sample, the resulting image of the mirror sphere gets distorted as it breaks the planar surface assumption.

6.4.2 Results

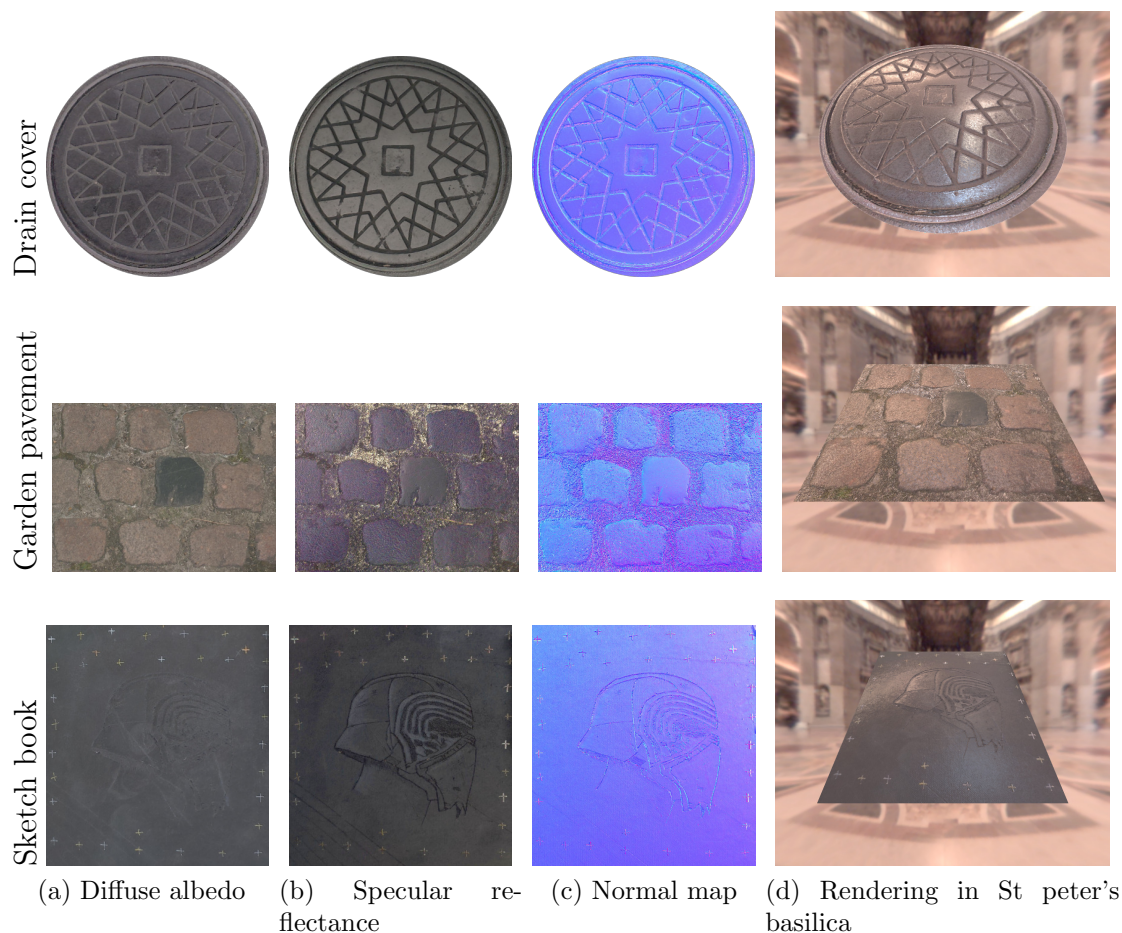


Figure 6.16: Reflectance maps ((a)-(c)) estimated using our mobile polarimetric setup. As we cannot estimate the specular roughness, we borrowed it from our DSLR setup for the “drain cover” and manually set it to 0.2 for the “garden pavement” and sketch book for rendering (d).

Figure 6.16 shows recovered reflectance maps, namely diffuse albedo (ρ_d), specular reflectance ($F(0^\circ)$) and surface normals (\vec{n}) for a few planar samples acquired with our mobile polarimetric setup. As can be seen, the maps of the “drain cover” (fig. 6.16, first row) are qualitatively very similar to those obtained with the DSLR setup (fig. 6.7, first row). Furthermore, even for the diffuse dominated “garden pavement” (fig. 6.16, second row), we are able to recover highly detailed surface normals and spatially-varying specular reflectance. Finally, the “sketch book” presents an interesting embossing pattern depicting Star Wars character “Kylo Ren” which is well captured by our normal and specular reflectance maps. Overall, although we are unable to recover specular roughness maps with this setup,

the other recovered maps already go a long way towards producing compelling renderings (fig. 6.16d).

We further validate our approach by comparing reflectance maps recovered from our mobile setup to both our DSLR setup (fig. 6.17) and the LCD-based method of Ghosh et al. [5] (fig. 6.18) and show good qualitative match in both cases, highlighting the practicality of the mobile polarimetric approach for on-site acquisition. While there are some visual differences, especially in the colour tone of the diffuse and specular reflectance maps estimated with our mobile setup, we attribute those to two main factors:

1. We were not able to deactivate the white balancing on the mobile phone, although this is somewhat mitigated by the radiometric calibration against the calibration target.
2. The transmission characteristics of the polarising filters employed for the DSLR camera (laminated glass) and mobile phone (laminated plastic) are different, with the former being of better quality than the latter.




Figure 6.17: Comparison of the reflectance maps obtained for the “drain cover” sample with our mobile polarimetric setup (top row) and DSLR setup (bottom row).




Figure 6.18: Comparison of the reflectance maps obtained for the “sketch book” sample with our mobile polarimetric setup and the LCD-based method of Ghosh et al. [5].

Outdoors Reflectometry - Summary

We have presented a novel approach for passive reflectometry for planar surfaces in completely uncontrolled outdoor environments using a combination of linear polarisation imaging, multi-view acquisition and inverse rendering. We demonstrated high quality estimation of spatially varying diffuse and specular reflectance, surface normals and specular roughness for a wide variety of planar real world materials ranging from diffuse dominated bricks and stone surfaces to very specular metal-dielectric composite surfaces. To the best of our knowledge, we are the first to apply polarisation based reflectometry in such complex and completely uncontrolled outdoor environments including busy urban settings. Unlike previous work on polarisation based shape/reflectance analysis which has assumed unpolarised or circularly polarised illumination, we take into account the potential partial linear polarisation of outdoor illumination and propose steps to mitigate the effects of such incident polarisation in our reflectance acquisition and analysis. While understandably not quite as accurate as completely controlled measurements, our method achieves sufficient accuracy for realistic rendering applications and is particularly suited for surfaces that are permanent on-site structures that cannot be brought indoors.

PART IV:

EPILOGUE

Chapter 7

Conclusion

In this dissertation, we have presented novel methods for on-site acquisition of surface reflectance based on off-the-shelf portable commodity hardware (part II). In particular, we have successfully demonstrated how an off-the-shelf mobile device can be utilised as a self-contained, economic active reflectometry device for measurements in semi-uncontrolled conditions indoors (chapters 3 and 4). We have further presented practical solutions to enhance the photo-realistic appearance (section 3.4) as well as transfer the reflectance properties (section 4.3) of the scanned materials by exploiting observations of the surfaces under natural illumination.

We have further proposed a novel passive reflectometric approach for permanent on-site surfaces found outdoors (part III), which we believe to be the first approach able to produce a complete set of highly detailed reflectance maps under uncontrolled natural illumination. As the latter is known to be partially linearly polarised in the most general case, we first derived a novel theoretical framework for polarisation imaging under the previously unstudied case of incident partial linear polarisation (chapter 5). Based on Mueller calculus, we have shown our derivations to be a generalisation of the well-studied theory of polarisation imaging assuming unpolarised incident illumination (section 5.2). Finally, from these derivations, we have developed a measurement protocol for passive reflectometry under natural illumination by noting that both theories (incident partially linearly polarised and unpolarised illumination) can be linked together under the right as-

assumptions (chapter 6). We have further proposed a practical hand-held acquisition setup for on-site mobile polarimetric reflectometry (section 6.4).

While the methods presented in this dissertation provide good qualitative and quantitative results that are comparable to other concurrent methods employing controlled or semi-controlled illumination setups, some limitations remain which we believe could give rise to interesting future research. First, an obvious future direction would be to extend our free-form mobile surface reflectometry approach to jointly recover reflectance and complete 3D geometry. Devices such as Google's Tango AR platform are a good example of what the future of mobile technology will look like, with powerful vision capabilities to enable hand-held free-form scanning. The recent work of Zhang et al. [101] is already going in that direction.

Another direction that we envision for future research is that of exploiting polarisation imaging to resolve anisotropy and to extend our derivations on polarisation imaging under partial linear polarisation to account for the complex index of refraction of metallic surfaces. It would also be interesting and straightforward to incorporate our "in-the-wild" polarimetric approach to the recent work of Cui et al. on polarimetric Multi-View Stereo (MVS) [102]. Furthermore, to make measurements more practical still, one could look at polarisation cameras where the common Bayer pattern has been replaced with micro-arrays of different polarisers to allow the acquisition of the complete Stokes field in a single shot [103]. In the same spirit, Kim et al. [104] have recently proposed a method for single-shot separation of layered reflectance, by inserting a pair of cross-polarised polarisers in the optical path of a light-field camera. While their method is currently limited to a single cross-polarisation pattern, adding a more complex pattern to allow complete Stokes measurements only seems natural.

It is our strong belief that the future of reflectometry is mobile and that as technology evolves, mobile phones and digital cameras will be able to go beyond the traditional 2D image pipeline to easily capture a true 3D representation of small as well as large scale structures, with which the user can interact.

“C’est pas faux.” - Perceval de Galles, dit Provençal le Gaulois

Bibliography

- [1] Jérémy Riviere, Pieter Peers, and Abhijeet Ghosh. Mobile surface reflectometry. *Computer Graphics Forum*, 35(1):191–202, 2016.
- [2] Jérémy Riviere, Ilya Reshetouski, Luka Filipi, and Abhijeet Ghosh. Polarization imaging reflectometry in the wild. *ACM Transaction on Graphics (TOG)*, 36(6):206, 2017.
- [3] JF Murray-Coleman and AM Smith. The automated measurement of brdfs and their application to luminaire modeling. *Journal of the Illuminating Engineering Society*, 19(1):87–99, 1990.
- [4] Paul Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of ACM SIGGRAPH 98*, 1998.
- [5] Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A. Wilson, and Paul E. Debevec. Estimating specular roughness and anisotropy from second order spherical gradient illumination. *Comput. Graph. Forum*, 28(4):1161–1170, 2009.
- [6] S. Shafer. Using color to separate reflection components. *COLOR research and Applications*, 10(4):210–218, 1985.
- [7] Fred E Nicodemus. Directional reflectance and emissivity of an opaque surface. *Applied optics*, 4(7):767–775, 1965.

- [8] Sing Choong Foo. *A gonioreflectometer for measuring the bidirectional reflectance of material for use in illumination computation*. PhD thesis, Cornell University, 1997.
- [9] Gregory J Ward. Measuring and modeling anisotropic reflection. *ACM SIGGRAPH Computer Graphics*, 26(2):265–272, 1992.
- [10] SR Marschner, SH Westin, EPF Lafortune, and KE Torrance. Image-based measurement of the bidirectional reflectance distribution function. *Appl. Opt.*, 39:2592–600, 2000.
- [11] Paul E Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. *ACM Transactions on Graphics (TOG)*, 1997.
- [12] Peiran Ren, Jiaping Wang, John Snyder, Xin Tong, and Baining Guo. Pocket reflectometry. *ACM Trans. Graph.*, 30(4):45:1–45:10, July 2011.
- [13] Vincent Masselus, Philip Dutré, and Frederik Anrys. The free-form light stage. In *ACM SIGGRAPH 2002 conference abstracts and applications*, pages 262–262. ACM, 2002.
- [14] Miika Aittala, Tim Weyrich, and Jaakko Lehtinen. Two-shot svbrdf capture for stationary materials. *ACM Transactions on Graphics*, 34(4):110, 2015.
- [15] Dar’ya Guarnera, Giuseppe Claudio Guarnera, Abhijeet Ghosh, Cornelia Denk, and Mashhuda Glencross. Brdf representation and acquisition. In *Computer Graphics Forum*, volume 35, pages 625–650. Wiley Online Library, 2016.
- [16] Giuseppe Claudio Guarnera, Abhijeet Ghosh, Ian Hall, Mashhuda Glencross, and Dar’ya Guarnera. Material capture and representation with applications in virtual reality. In *ACM SIGGRAPH 2017 Courses*, page 6. ACM, 2017.
- [17] Peter S Shirley. *Physically based lighting calculations for computer graphics*. PhD thesis, University of Illinois at Urbana-Champaign, 1991.

- [18] Robert L Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics (TOG)*, 1(1):7–24, 1982.
- [19] Christophe Schlick. An inexpensive brdf model for physically-based rendering. In *Computer graphics forum*, volume 13, pages 233–246. Wiley Online Library, 1994.
- [20] Bui Tuong-Phong. Illumination for computer-generated images. *University of Utah*, pages 29–51, 1973.
- [21] Eric P Lafortune and Yves D Willems. Using the modified phong reflectance model for physically based rendering. 1994.
- [22] Christian Schüler. An efficient and physically plausible real time shading model. *ShaderX7, Charles River Media*, pages 175–187, 2009.
- [23] James F Blinn. Models of light reflection for computer synthesized pictures. In *ACM SIGGRAPH Computer Graphics*, volume 11, pages 192–198. ACM, 1977.
- [24] Fabian Giesen. Phong normalization factor derivation. <http://www.farbrausch.de/~fg/stuff/phong.pdf>.
- [25] Michael Ashikhmin and Peter Shirley. An anisotropic phong brdf model. *Journal of graphics tools*, 5(2):25–32, 2000.
- [26] Kenneth E Torrance and Ephraim M Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57(9):1105–1114, 1967.
- [27] Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*, pages 195–206. Eurographics Association, 2007.
- [28] Jos Stam. An illumination model for a skin layer bounded by rough surfaces. In *Rendering Techniques 2001*, pages 39–52. Springer, 2001.

- [29] André Spizzichino and Petr Beckmann. *The Scattering of Electromagnetic Waves from Rough Surfaces*. New York, Paris, 1963.
- [30] James Frederick Blinn. Computer display of curved surfaces. 1978.
- [31] B Smith. Geometrical shadowing of a random rough surface. *IEEE transactions on antennas and propagation*, 15(5):668–671, 1967.
- [32] G Brown. Shadowing by non-gaussian random surfaces. *IEEE Transactions on Antennas and Propagation*, 28(6):788–790, 1980.
- [33] Christophe Bourlier, Gérard Berginc, and Joseph Saillard. One-and two-dimensional shadowing functions for any height and slope stationary uncorrelated surface in the monostatic and bistatic configurations. *IEEE Transactions on Antennas and Propagation*, 50(3):312–324, 2002.
- [34] TS Trowbridge and KP Reitz. Average irregularity representation of a rough surface for ray reflection. *JOSA*, 65(5):531–536, 1975.
- [35] Michael Oren and Shree K Nayar. Generalization of lambert’s reflectance model. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 239–246. ACM, 1994.
- [36] CURET. Curet: Columbia-utrech reflectance and texture,. <http://www.cs.columbia.edu/CAVE/curet/>, 1999.
- [37] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. A data-driven reflectance model. *ACM Transactions on Graphics (TOG)*, 22(3):759–769, 2003.
- [38] Andrew Gardner, Chris Tchou, Tim Hawkins, and Paul Debevec. Linear light source reflectometry. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 22(3):749–758, 2003.
- [39] Guojun Chen, Yue Dong, Pieter Peers, Jiawan Zhang, and Xin Tong. Reflectance scanning: Estimating shading frame and brdf with generalized linear light sources. *ACM Trans. Graph.*, 33(4):117:1–117:11, July 2014.

- [40] Yannick Francken, Tom Cuypers, Tom Mertens, and Philippe Bekaert. Gloss and normal map acquisition of mesostructures using gray codes. In *Proceedings of the 5th International Symposium on Advances in Visual Computing: Part II, ISVC '09*, pages 788–798, Berlin, Heidelberg, 2009. Springer-Verlag.
- [41] Chun-Po Wang, Noah Snavely, and Steve Marschner. Estimating dual-scale properties of glossy surfaces from step-edge lighting. *ACM Trans. Graph.*, 30(6):172:1–172:12, December 2011.
- [42] Miika Aittala, Tim Weyrich, and Jaakko Lehtinen. Practical SVBRDF capture in the frequency domain. *ACM Trans. on Graphics (Proc. SIGGRAPH)*, 32(4), 2013.
- [43] Hongzhi Wu and Kun Zhou. Appfusion: Interactive appearance acquisition using a kinect sensor. *Computer Graphics Forum*, 34(6):289–298, 2015.
- [44] Nicolas Tisserand and Burrus Nicolas. Skanect. <http://skanect.occipital.com/>, 2011.
- [45] ProFactor. Reconstructme. <http://reconstructme.net/>.
- [46] Miika Aittala, Timo Aila, and Jaakko Lehtinen. Reflectance modeling by neural texture synthesis. *ACM Transactions on Graphics*, 35(4), 2016.
- [47] Xiao Li, Yue Dong, Pieter Peers, and Xin Tong. Modeling surface appearance from a single photograph using self-augmented convolutional neural networks. *ACM Transactions on Graphics (TOG)*, 36(4):45, 2017.
- [48] Mashhuda Glencross, Gregory J Ward, Francho Melendez, Caroline Jay, Jun Liu, and Roger Hubbard. A perceptually validated model for surface depth hallucination. *ACM Transactions on Graphics (TOG)*, 27(3):59, 2008.
- [49] Fabiano Romeiro, Yuriy Vasilyev, and Todd Zickler. Passive reflectometry. In *Computer Vision—ECCV 2008*, pages 859–872. Springer, 2008.
- [50] Fabiano Romeiro and Todd Zickler. Blind reflectometry. In *Proc. of ECCV*, pages 45–58, 2010.

- [51] Stephen Lombardi and Ko Nishino. Reflectance and natural illumination from a single image. In *Proc. European Conference on Computer Vision (ECCV'12)*, pages 582–595, 2012.
- [52] Geoffrey Oxholm and Ko Nishino. Shape and reflectance from natural illumination. In *Proc. European Conference on Computer Vision (ECCV'12)*, pages 528–541, 2012.
- [53] Daniel Hauagge, Scott Wehrwein, Paul Upchurch, Kavita Bala, and Noah Snavely. Reasoning about photo collections using models of outdoor illumination. In *Proceedings of the British Machine Vision Conference. BMVA Press*, 2014.
- [54] Yue Dong, Guojun Chen, Pieter Peers, Jiawan Zhang, and Xin Tong. Appearance-from-motion: recovering spatially varying surface reflectance under unknown lighting. *ACM Transactions on Graphics (TOG)*, 33(6):193, 2014.
- [55] Rui Xia, Yue Dong, Pieter Peers, and Xin Tong. Recovering shape and spatially-varying surface reflectance under unknown illumination. *ACM Trans. Graph.*, 35(6):187:1–187:12, November 2016.
- [56] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar. Acquiring the reflectance field of a human face. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 145–156. ACM Press/Addison-Wesley Publishing Co., 2000.
- [57] Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, and Paul Debevec. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Rendering Techniques*, pages 183–194, 2007.

- [58] Abhijeet Ghosh, Tim Hawkins, Pieter Peers, Sune Frederiksen, and Paul Debevec. Practical modeling and acquisition of layered facial reflectance. *ACM Transactions on Graphics*, 27(5):139:1–139:10, December 2008.
- [59] Lawrence B. Wolff and Terrance E. Boult. Constraining object features using a polarization reflectance model. *PAMI*, 13(7):635–657, 1991.
- [60] Volker Müller. Polarization-based separation of diffuse and specular surface-reflection. In *Mustererkennung 1995*, pages 202–209. Springer, 1995.
- [61] Kazutada Koshikawa. A polarimetric approach to shape understanding of glossy objects. pages 190–192, 1992.
- [62] Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A. Wilson, and Paul Debevec. Circularly polarized spherical illumination reflectometry. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)*, 29:162:1–162:12, December 2010.
- [63] L. B. Wolff. Polarization-based material classification from specular reflection. *PAMI*, 12(11):1059–1071, 1990.
- [64] Gary Atkinson and Edwin R Hancock. Recovery of surface orientation from diffuse polarization. *Image Processing, IEEE Transactions on*, 15(6):1653–1664, 2006.
- [65] Abhijeet Ghosh, Graham Fyffe, Borom Tunwattanapong, Jay Busch, Xueming Yu, and Paul Debevec. Multiview face capture using polarized spherical gradient illumination. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, pages 129:1–129:10, 2011.
- [66] Daisuke Miyazaki, Robby T. Tan, Kenji Hara, and Katsushi Ikeuchi. Polarization-based inverse rendering from a single view. In *ICCV*, pages 982–987, 2003.
- [67] L. B. Wolff. Surface orientation from two camera stereo with polarizers. In *Proc. SPIE Conf. Optics, Illumination and Image Sensing for Machine Vision IV*, volume 1194, pages 287–297, 1989.

- [68] Megumi Saito, Yoichi Sato, Katsushi Ikeuchi, and Hiroshi Kashiwagi. Measurement of surface orientations of transparent objects by use of polarization in highlight. *J. Opt. Soc. Am. A*, 16(9):2286–2293, 1999.
- [69] Daisuke Miyazaki, Masataka Kagesawa, and Katsushi Ikeuchi. Polarization-based transparent surface modeling from two views. In *ICCV*, pages 1381–1386, 2003.
- [70] Firooz Sadjadiz and Farzad Sadjadi. Extraction of surface normal and index of refraction using a pair of passive infrared polarimetric sensors. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–5, 2007.
- [71] Giuseppe Claudio Guarnera, Pieter Peers, Paul E. Debevec, and Abhijeet Ghosh. Estimating surface normals from spherical stokes reflectance fields. In *ECCV Workshop on Color and Photometry in Computer Vision*, pages 340–349, 2012.
- [72] Achuta Kadambi, Vage Taamazyan, Boxin Shi, and Ramesh Raskar. Polarized 3d: High-quality depth sensing with polarization cues. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3370–3378, 2015.
- [73] W.A.P. Smith, R. Ramamoorthi, and S. Tozza. Linear depth estimation from an uncalibrated, monocular polarisation image. In *Proc. of European Conference on Computer Vision (ECCV)*, pages 517–526, 2016.
- [74] MobileFuture. The rise of mobile: 11.6 billion mobile-connected devices by 2020. <http://mobilefuture.org/the-rise-of-mobile-11-6-billion-mobile-connected-devices-by-2020>. Online: 04/02/2016.
- [75] Michael Ashikhmin and Simon Premože. Distribution-based BRDFs, 2007. <http://www.cs.utah.edu/premoze/dbrdf/>.
- [76] Chris Harris and Mike Stephens. A combined corner and edge detector. In *Proc. of Fourth Alvey Vision Conference*, pages 147–151, 1988.

- [77] Jean-Yves Bouguet. Pyramidal implementation of the affine lucas kanade feature tracker description of the algorithm. *Intel Corporation*, 5, 2001.
- [78] Georg Klein and David Murray. Parallel tracking and mapping for small ar workspaces. In *Mixed and Augmented Reality, 2007. ISMAR 2007. 6th IEEE and ACM International Symposium on*, pages 225–234. IEEE, 2007.
- [79] Abe Davis, Marc Levoy, and Fredo Durand. Unstructured light fields. In *Computer Graphics Forum*, volume 31, pages 305–314. Wiley Online Library, 2012.
- [80] Jan Jachnik, Richard A Newcombe, and Andrew J Davison. Real-time surface light-field capture for augmentation of planar specular surfaces. In *Mixed and Augmented Reality (ISMAR), 2012 IEEE International Symposium on*, pages 91–97. IEEE, 2012.
- [81] Mark A Robertson, Sean Borman, and Robert L Stevenson. Dynamic range improvement through multiple exposures. In *Image Processing, 1999. ICIP 99. Proceedings. 1999 International Conference on*, volume 3, pages 159–163. IEEE, 1999.
- [82] Thabo Beeler, Bernd Bickel, Paul Beardsley, Bob Sumner, and Markus Gross. High-quality single-shot capture of facial geometry. *ACM Trans. Graph.*, 29:40:1–40:9, July 2010.
- [83] John William Strutt. Lviii. on the scattering of light by small particles. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 41(275):447–454, 1871.
- [84] John W Strutt. Xv. on the light from the sky, its polarization and colour. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 41(271):107–120, 1871.
- [85] Daisuke Miyazaki and Katsushi Ikeuchi. Inverse polarization raytracing: Estimating surface shapes of transparent objects. In *CVPR*, pages 910–917, 2005.

- [86] E. Collett. *Field Guide to Polarization, SPIE Field Guides vol. FG05*. SPIE, 2005.
- [87] Hua Chen and Lawrence B Wolff. Polarization phase-based method for material classification in computer vision. *International Journal of Computer Vision*, 28(1):73–83, 1998.
- [88] Daisuke Miyazaki, Takuya Shigetomi, Masashi Baba, Ryo Furukawa, Shinsaku Hiura, and Naoki Asada. Polarization-based surface normal estimation of black specular objects from multiple viewpoints. In *3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2012 Second International Conference on*, pages 104–111. IEEE, 2012.
- [89] Cong Phuoc Huynh, Antonio Robles-Kelly, and Edwin Hancock. Shape and refractive index recovery from single-view polarisation images. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1229–1236. IEEE, 2010.
- [90] G.P. Können. *Polarized Light in Nature*. Cambridge University Press, 1985.
- [91] Jannik Boll Nielsen, Henrik Wann Jensen, and Ravi Ramamoorthi. On optimal, minimal brdf sampling for reflectance acquisition. *ACM Trans. Graph.*, 34(6):186:1–186:11, October 2015.
- [92] Rafal Mantiuk, Grzegorz Krawczyk, Radoslaw Mantiuk, and Hans-Peter Seidel. High dynamic range imaging pipeline: Perception-motivated representation of visual content. In *Electronic Imaging 2007*, pages 649212–649212. International Society for Optics and Photonics, 2007.
- [93] Changchang Wu. *Visualsfm: A visual structure from motion system*. 2011.
- [94] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.

- [95] Changchang Wu, Sameer Agarwal, Brian Curless, and Steven M Seitz. Multicore bundle adjustment. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 3057–3064. IEEE, 2011.
- [96] Chloe LeGendre, Xueming Yu, Dai Liu, Jay Busch, Andrew Jones, Sumanta Pattanaik, and Paul Debevec. Practical multispectral lighting reproduction. *ACM Transactions on Graphics (TOG)*, 35(4):32, 2016.
- [97] Pixel and Poly. Pixelandpoly ior list. <https://pixelandpoly.com/ior.html>.
- [98] CGSociety. A complete ior list. <http://forums.cgsociety.org/showthread.php?t=513458>, 2007.
- [99] Craig Donner and Henrik Wann Jensen. A spectral bssrdf for shading human skin. *Rendering techniques*, 2006:409–418, 2006.
- [100] KE Peiponen, V Kontturi, I Niskanen, M Juuti, J Rätty, H Koivula, and M Toivakka. On estimation of complex refractive index and colour of dry black and cyan offset inks by a multi-function spectrophotometer. *Measurement Science and Technology*, 19(11):115601, 2008.
- [101] Edward Zhang, Michael F Cohen, and Brian Curless. Emptying, refurbishing, and relighting indoor spaces. *ACM Transactions on Graphics (TOG)*, 35(6):174, 2016.
- [102] Zhaopeng Cui, Jinwei Gu, Boxin Shi, Ping Tan, and Jan Kautz. Polarimetric multi-view stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1558–1567, 2017.
- [103] Ricoh. Polarisation camera.
- [104] Jaewon Kim, Shahram Izadi, and Abhijeet Ghosh. Single-shot layered reflectance separation using a polarized light field camera. 2016.


Appendices

A Additional renderings - Mobile Surface Reflectometry (part II)


More rendering results are provided as animations in the accompanying video, available at: <https://goo.gl/cZGM1x>¹



(a) Eucalyptus grove



(b) Grace cathedral



(c) St peter's basilica



(d) Uffizi gallery

Figure 1: “Anniversary” greeting card: Environmental illumination renderings.

¹Please consider downloading the video in case the video playback is of poor quality.



(a) Eucalyptus grove



(b) Grace cathedral




(c) St peter's basilica




(d) Uffizi gallery


Figure 2: "New job" greeting card: Environmental illumination renderings.



(a) Eucalyptus grove



(b) Grace cathedral




(c) St peter's basilica



(d) Uffizi gallery

Figure 3: “Twenty-one” greeting card: Environmental illumination renderings.



(a) Eucalyptus grove



(b) Grace cathedral



(c) St peter's basilica



(d) Uffizi gallery

Figure 4: **“Bronze medal”**: Environmental illumination renderings.



(a) Eucalyptus grove



(b) Grace cathedral



(c) St peter's basilica



(d) Uffizi gallery


Figure 5: “Cha” tea box cover: Environmental illumination renderings.



(a) Eucalyptus grove



(b) Grace cathedral



(c) St peter's basilica



(d) Uffizi gallery

Figure 6: “Child’s book” cover: Environmental illumination renderings.




(a) Eucalyptus grove



(b) Grace cathedral



(c) St Peter's basilica



(d) Uffizi gallery

Figure 7: “Intel” coin: Environmental illumination renderings.

B Additional renderings - Outdoors reflectometry (part III)

More rendering results are provided as animations in the accompanying video, available at: www.google.com



(a) Eucalyptus grove



(b) Grace cathedral




(c) St peter's basilica




(d) Uffizi gallery


Figure 8: “**Drain cover**”: Environmental illumination renderings.




(a) Eucalyptus grove



(b) Grace cathedral




(c) St peter's basilica




(d) Uffizi gallery


Figure 9: “Red book” cover: Environmental illumination renderings.



(a) Eucalyptus grove



(b) Grace cathedral




(c) St peter's basilica




(d) Uffizi gallery


Figure 10: “Red bricks”: Environmental illumination renderings.



(a) Eucalyptus grove



(b) Grace cathedral




(c) St peter's basilica




(d) Uffizi gallery


Figure 11: “**Garden pavement**”: Environmental illumination renderings.




(a) Eucalyptus grove



(b) Grace cathedral



(c) St peter's basilica




(d) Uffizi gallery

Figure 12: “**Stone pavement**”: Environmental illumination renderings.



(a) Eucalyptus grove



(b) Grace cathedral




(c) St Peter's basilica




(d) Uffizi gallery


Figure 13: “**Bas-relief**”: Environmental illumination renderings.



(a) Eucalyptus grove



(b) Grace cathedral



(c) St peter's basilica



(d) Uffizi gallery

Figure 14: “Sketch book” cover: Environmental illumination renderings.

C Copyright notices

**JOHN WILEY AND SONS LICENSE
TERMS AND CONDITIONS**

Jan 30, 2017

This Agreement between J r my Riviere ("You") and John Wiley and Sons ("John Wiley and Sons") consists of your license details and the terms and conditions provided by John Wiley and Sons and Copyright Clearance Center.

License Number	4038850051065
License date	Jan 30, 2017
Licensed Content Publisher	John Wiley and Sons
Licensed Content Publication	Computer Graphics Forum
Licensed Content Title	Mobile Surface Reflectometry
Licensed Content Author	J. Riviere,P. Peers,A. Ghosh
Licensed Content Date	Aug 25, 2015
Licensed Content Pages	12
Type of use	Dissertation/Thesis
Requestor type	Author of this Wiley article
Format	Print and electronic
Portion	Full article
Will you be translating?	No
Title of your thesis / dissertation	On-site acquisition of surface reflectance
Expected completion date	Aug 2017
Expected size (number of pages)	150
Requestor Location	J�r�my Riviere 101 winders Road London, SW113HD United Kingdom Attn: J�r�my Riviere
Publisher Tax ID	EU826007151
Billing Type	Invoice
Billing Address	J�r�my Riviere 101 winders Road London, United Kingdom SW113HD Attn: J�r�my Riviere
Total	0.00 GBP
Terms and Conditions	

TERMS AND CONDITIONS

This copyrighted material is owned by or exclusively licensed to John Wiley & Sons, Inc. or one of its group companies (each a "Wiley Company") or handled on behalf of a society with which a Wiley Company has exclusive publishing rights in relation to a particular work (collectively "WILEY"). By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the billing and payment terms and conditions established by the Copyright Clearance Center Inc., ("CCC's Billing and Payment terms and conditions"), at the time that you opened your RightsLink account (these are available at any time at <http://myaccount.copyright.com>).

Terms and Conditions

- The materials you have requested permission to reproduce or reuse (the "Wiley Materials") are protected by copyright.
- You are hereby granted a personal, non-exclusive, non-sub licensable (on a stand-alone basis), non-transferable, worldwide, limited license to reproduce the Wiley Materials for the purpose specified in the licensing process. This license, **and any CONTENT (PDF or image file) purchased as part of your order**, is for a one-time use only and limited to any maximum distribution number specified in the license. The first instance of republication or reuse granted by this license must be completed within two years of the date of the grant of this license (although copies prepared before the end date may be distributed thereafter). The Wiley Materials shall not be used in any other manner or for any other purpose, beyond what is granted in the license. Permission is granted subject to an appropriate acknowledgement given to the author, title of the material/book/journal and the publisher. You shall also duplicate the copyright notice that appears in the Wiley publication in your use of the Wiley Material. Permission is also granted on the understanding that nowhere in the text is a previously published source acknowledged for all or part of this Wiley Material. Any third party content is expressly excluded from this permission.
- With respect to the Wiley Materials, all rights are reserved. Except as expressly granted by the terms of the license, no part of the Wiley Materials may be copied, modified, adapted (except for minor reformatting required by the new Publication), translated, reproduced, transferred or distributed, in any form or by any means, and no derivative works may be made based on the Wiley Materials without the prior permission of the respective copyright owner. **For STM Signatory Publishers clearing permission under the terms of the [STM Permissions Guidelines](#) only, the terms of the license are extended to include subsequent editions and for editions in other languages, provided such editions are for the work as a whole in situ and does not involve the separate exploitation of the permitted figures or extracts**, You may not alter, remove or suppress in any manner any copyright, trademark or other notices displayed by the Wiley Materials. You may not license, rent, sell, loan, lease, pledge, offer as security, transfer or assign the Wiley Materials on a stand-alone basis, or any of the rights granted to you hereunder to any other person.
- The Wiley Materials and all of the intellectual property rights therein shall at all times remain the exclusive property of John Wiley & Sons Inc, the Wiley Companies, or their respective licensors, and your interest therein is only that of having possession of and the right to reproduce the Wiley Materials pursuant to Section 2 herein during the continuance of this Agreement. You agree that you own no right, title or interest in or to the Wiley Materials or any of the intellectual property rights therein. You shall have no rights hereunder other than the license as provided for above in Section 2. No right, license or interest to any trademark, trade name, service mark or other branding ("Marks") of WILEY or its licensors is granted hereunder, and you agree that you shall not assert any such right, license or interest with respect thereto
- NEITHER WILEY NOR ITS LICENSORS MAKES ANY WARRANTY OR REPRESENTATION OF ANY KIND TO YOU OR ANY THIRD PARTY, EXPRESS, IMPLIED OR STATUTORY, WITH RESPECT TO THE MATERIALS OR THE ACCURACY OF ANY INFORMATION CONTAINED IN THE MATERIALS, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTY OF MERCHANTABILITY, ACCURACY, SATISFACTORY QUALITY, FITNESS FOR A PARTICULAR PURPOSE, USABILITY, INTEGRATION OR NON-INFRINGEMENT AND ALL SUCH WARRANTIES ARE HEREBY EXCLUDED BY WILEY AND ITS LICENSORS AND WAIVED BY YOU.
- WILEY shall have the right to terminate this Agreement immediately upon breach of this Agreement by you.

- You shall indemnify, defend and hold harmless WILEY, its Licensors and their respective directors, officers, agents and employees, from and against any actual or threatened claims, demands, causes of action or proceedings arising from any breach of this Agreement by you.
- IN NO EVENT SHALL WILEY OR ITS LICENSORS BE LIABLE TO YOU OR ANY OTHER PARTY OR ANY OTHER PERSON OR ENTITY FOR ANY SPECIAL, CONSEQUENTIAL, INCIDENTAL, INDIRECT, EXEMPLARY OR PUNITIVE DAMAGES, HOWEVER CAUSED, ARISING OUT OF OR IN CONNECTION WITH THE DOWNLOADING, PROVISIONING, VIEWING OR USE OF THE MATERIALS REGARDLESS OF THE FORM OF ACTION, WHETHER FOR BREACH OF CONTRACT, BREACH OF WARRANTY, TORT, NEGLIGENCE, INFRINGEMENT OR OTHERWISE (INCLUDING, WITHOUT LIMITATION, DAMAGES BASED ON LOSS OF PROFITS, DATA, FILES, USE, BUSINESS OPPORTUNITY OR CLAIMS OF THIRD PARTIES), AND WHETHER OR NOT THE PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY PROVIDED HEREIN.
- Should any provision of this Agreement be held by a court of competent jurisdiction to be illegal, invalid, or unenforceable, that provision shall be deemed amended to achieve as nearly as possible the same economic effect as the original provision, and the legality, validity and enforceability of the remaining provisions of this Agreement shall not be affected or impaired thereby.
- The failure of either party to enforce any term or condition of this Agreement shall not constitute a waiver of either party's right to enforce each and every term and condition of this Agreement. No breach under this agreement shall be deemed waived or excused by either party unless such waiver or consent is in writing signed by the party granting such waiver or consent. The waiver by or consent of a party to a breach of any provision of this Agreement shall not operate or be construed as a waiver of or consent to any other or subsequent breach by such other party.
- This Agreement may not be assigned (including by operation of law or otherwise) by you without WILEY's prior written consent.
- Any fee required for this permission shall be non-refundable after thirty (30) days from receipt by the CCC.
- These terms and conditions together with CCC's Billing and Payment terms and conditions (which are incorporated herein) form the entire agreement between you and WILEY concerning this licensing transaction and (in the absence of fraud) supersedes all prior agreements and representations of the parties, oral or written. This Agreement may not be amended except in writing signed by both parties. This Agreement shall be binding upon and inure to the benefit of the parties' successors, legal representatives, and authorized assigns.
- In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall prevail.
- WILEY expressly reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.
- This Agreement will be void if the Type of Use, Format, Circulation, or Requestor Type was misrepresented during the licensing process.

- This Agreement shall be governed by and construed in accordance with the laws of the State of New York, USA, without regards to such state's conflict of law rules. Any legal action, suit or proceeding arising out of or relating to these Terms and Conditions or the breach thereof shall be instituted in a court of competent jurisdiction in New York County in the State of New York in the United States of America and each party hereby consents and submits to the personal jurisdiction of such court, waives any objection to venue in such court and consents to service of process by registered or certified mail, return receipt requested, at the last known address of such party.

WILEY OPEN ACCESS TERMS AND CONDITIONS

Wiley Publishes Open Access Articles in fully Open Access Journals and in Subscription journals offering Online Open. Although most of the fully Open Access journals publish open access articles under the terms of the Creative Commons Attribution (CC BY) License only, the subscription journals and a few of the Open Access Journals offer a choice of Creative Commons Licenses. The license type is clearly identified on the article.

The Creative Commons Attribution License

The [Creative Commons Attribution License \(CC-BY\)](#) allows users to copy, distribute and transmit an article, adapt the article and make commercial use of the article. The CC-BY license permits commercial and non-

Creative Commons Attribution Non-Commercial License

The [Creative Commons Attribution Non-Commercial \(CC-BY-NC\) License](#) permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.(see below)

Creative Commons Attribution-Non-Commercial-NoDerivs License

The [Creative Commons Attribution Non-Commercial-NoDerivs License \(CC-BY-NC-ND\)](#) permits use, distribution and reproduction in any medium, provided the original work is properly cited, is not used for commercial purposes and no modifications or adaptations are made. (see below)

Use by commercial "for-profit" organizations

Use of Wiley Open Access articles for commercial, promotional, or marketing purposes requires further explicit permission from Wiley and will be subject to a fee.

Further details can be found on Wiley Online Library

<http://olabout.wiley.com/WileyCDA/Section/id-410895.html>

Other Terms and Conditions:

v1.10 Last updated September 2015

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.