

Precise and accurate image-based measurements of optical material properties

Présentée le 23 octobre 2020

à la Faculté informatique et communications
Laboratoire d'images et représentation visuelle
Programme doctoral en informatique et communications

pour l'obtention du grade de Docteur ès Sciences

par

Marjan SHAHPASKI

Acceptée sur proposition du jury

Prof. M. Pauly, président du jury
Prof. S. Süsstrunk, directrice de thèse
Dr M. Hébert, rapporteur
Dr L. R. Sapaico Valera, rapporteur
Prof. R. D. Hersch, rapporteur

Those who fall in love with practice without science are like a sailor who enters a ship
without a helm or a compass, and who never can be certain whither he is going.
— Leonardo da Vinci

To my family...

Acknowledgements

This marks the end of my biggest project thus far. It was a very formative experience on an academic, professional and personal level. It was a period with many ups and downs, good and bad, beauty and terror. It took six years, thus one could expect nothing less. Here I would like to acknowledge the people who rode the waves with me, who helped me just keep going, and turn the downs into ups, the bad into good, the terror into beauty, since no feeling is final.

First of all, I would like to thank my thesis advisor, Prof. Sabine Süsstrunk. I am very fortunate to have gotten the chance to work on this PhD with her, and I am very grateful for her availability, scientific advice, freedom and support throughout the whole process. Furthermore, her patience was invaluable, especially around half way into the thesis, when I was chronically tired and sleep deprived — I cannot believe that any other advisor would have had so much understanding. Our meetings were always intensely motivating for me, which helped me persevere and reach the end. Thank you Sabine for never doubting me, and always being there for me. You are more than a mentor, you are a true friend.

I would like to express my gratitude to the members of my thesis committee, Prof. Mark Pauly, Prof. Roger D. Hersch, Dr. Mathieu Hébert and Dr. Luis Ricardo Sapaico. It was an honor and delight to discuss my work with them. Their constructive feedback certainly made this thesis better, especially Ricardo's, with whom we had numerous meetings and exchanges. I am also thankful to Prof. Wenzel Jakob for sharing advice on BTDF measurement.

Being part of the Image and Visual Representation Lab (IVRL) comes with real privileges. I was lucky to share the office with Edo Collins, with whom we had numerous engaging discussions related to our research findings, but also on global, cantonal, municipal, and personal levels. We might have graduated a year sooner without them, but it was indeed time well spent. I had the pleasure to collaborate on a project with Majed El Helou, and to successfully compete in a computer vision challenge alongside him, Ruofan Zhou, and Fayez Lahoud. I would also like to thank the past and current members of the lab: Adrien, Albrecht, Bahar, Baran, Bin, Chen, Damian, Deblina, Gökhan, Hakgu, Hakki, Joachim, Leo, Mahdi, Martijn, Nick, RK, Sami, Seungryong, Siavash, Zahra, with whom I was fortunate to engage in various discussions, to take part in lab outings, and numerous celebrations of their successes. A great thank you goes to Françoise Behn, the administrative assistant of the lab, for helping me navigate the bureaucratic maze.

Acknowledgements

Continuing with the academic circle, I was fortunate to share the floor with brilliant people from our neighboring labs. Starting with my dear colleagues from LCAV, I enjoyed the conversations I had with Frederike and Sepand during our coffee breaks and over lunches, but also the company of the always positive Mihailo, Karen, Miranda, Gilles, Adam, Adrien, Michalina. I would be remiss not to thank my colleagues from the computer graphics labs, namely Mina, with whom I started the PhD, and Tizian, Merlin and Delio for helping me modify Mitsuba. Also, I spent an amazing summer with Morgane, Tom, and Petar, under the guidance of Dr. Anselm Grundhöfer during my internship at Disney Research Zürich.

Having people that speak your native language is a great advantage, because they can sense the nuances in the speech, which makes the interactions much richer. Ana, Darko, Dino, Elena, Elizabeta, Emil, Gorica, Igor, Kristina, Mladen, Nikola, Nikolche, Vase, and many others. I thank them for all the birthday parties, dinners, board-game nights, and great moments in general, which should have been much more numerous!

I was also lucky to have the support from friends and family scattered in various places around the world. I am very grateful to my cousin Ratko for welcoming me, and helping me out from my first day in Switzerland. I would also like to thank my very old and close friends Martin, Ivana and Goran for their support, positive attitude and numerous excursions within and outside of Switzerland. I am very lucky to have Trajko as my oldest friend, to whom I am indebted for the moral support and unbreakable spirit, and with whom one can never feel down for too long.

I was very fortunate to be accompanied on this journey by my wife Maja. We were together during all the ups and downs, good and bad, beauty and terror along the ride. We visited some amazing places, and we were extremely fortunate to have our son Maxim. I am especially grateful for the period around the oral thesis exam, where you took over the complete care for our son. He is incredibly sweet, but also a formidable explorer that stops at nothing. Thank you for filling our home with laughter, and our hearts with joy and happiness. Without you, this endeavor would not have been half as interesting.

Finally, I would like to thank my parents, Mileva and Branko, for their selfless and unwavering support throughout my whole life. They were always there for me, provided me with an unbelievable childhood, good education, and excellent conditions in which I could thrive. Their life guidance instilled good core values in me, taught me to never settle, and continually pursue what I love. Furthermore, they helped me directly to bring this work to a close by taking care of Maxim while I was writing this thesis. I also enjoyed the full support and unconditional love of my sister Maja, who is absolutely the best sister one could ask for. Accompanied by my brother-in-law Dean, and my adorable niece Marta, they make for a fantastic team. You have set the foundations for this journey to begin, and your support has helped me reach the finish line. For that, I will be eternally grateful.

Lausanne, September 28, 2020

Marjan Shahpaski

Abstract

The appearance of objects is governed by how they reflect, transmit and absorb the light they receive. That, in turn, depends on the material's internal structure, surface structure, and viewing and illumination directions. Changes in those characteristics can produce dramatic shifts in a material's appearance.

In this thesis, we explore how the optical properties can be estimated accurately using computational imaging. We focus on those properties that have an impact on material appearance. Current solutions are based on using instruments that are on the expensive side, such as spectrophotometers, profilometers and goniophotometers; hence, we explore the use of more affordable equipment, such as cameras. We start with the reflected light, and how the surface structure and roughness can be measured with commodity hardware such as projectors and cameras. We continue to the transmitted light, where we study methods for measuring hemispherical transmittance, and then go one step further and propose a novel design for the measurement of directional transmittance.

Systems based on the projection of structured light can be used for high-precision depth estimation. These methods provide non-contact means of depth estimation by using a calibrated camera-projector pair. Calibration of such systems is in general a cumbersome and time-consuming task. We, thus, present a novel method that allows for simultaneous geometric and radiometric calibration of a projector-camera pair. It is automatic, simple, and does not require specialized hardware. We prewarp and align a specially designed projection pattern onto a printed pattern of different colorimetric properties. After capturing the patterns in several orientations, we perform geometric calibration by estimating the corner locations of the two patterns in different color channels. We perform radiometric calibration of the projector by using the information contained inside the projected squares. We show that our method performs on par with current approaches that all require separate geometric and radiometric calibration. This makes it efficient, user friendly, and faster compared to dedicated geometric and radiometric calibration techniques. We then apply the algorithm to calibrate a structured light system with the goal of surface structure estimation of nearly planar surfaces. The surface structure is then used to calculate the macroscopic roughness. Over a set of samples, our results are comparable to those from a high-grade commercial profilometer.

The accurate measurement of transmittance properties of materials is essential in the printing and display industries in order to ensure accurate color reproduction. We explore different

Abstract

measurement geometries for total transmittance, and show that the transmittance measurements are highly affected by the geometry used, since certain geometries can introduce a measurement bias. We build a flexible custom setup that can simulate these geometries, and evaluate it both qualitatively and quantitatively over a set of samples with varied optical properties. We also compare our measurements against those of widely used commercial solutions, and show that significant differences exist over our test set. However, when the bias is correctly compensated, we observe very low differences. These findings, therefore, stress the importance of including the measurement geometry when reporting total transmittance. The bidirectional scattering distribution function (BSDF) is of great importance in many applications, from rendering images and visual special effects to architectural modelling of illumination. We propose a novel design for the measurement of BSDF, more specifically the bidirectional transmittance distribution function (BTDF). Our approach is based on radial catadioptric imaging, where we illuminate the measured sample with a collimated beam of light, the emerging light on the other side of the sample is then collected by a diffusing cylinder, and, after many internal reflections, a camera registers the steady-state radiance of the cylinder's interior. After inverting the light transport inside the cylinder, we can compute the sample's BTDF. Our approach effectively trades mechanical complexity with computational complexity. In simulations, our proposed design achieves results that are well aligned with the ground truth, for both smooth and high-frequency BTDFs that feature sharp intensity changes.

Keywords: computational imaging, structured light, depth estimation, transmittance, BRDF, BTDF, light transport, goniophotometer, radial imaging, catadioptric imaging.

Résumé

L'apparence des objets est régie par la façon dont ils réfléchissent, transmettent et absorbent la lumière qu'ils reçoivent. Cela dépend de la structure interne du matériau, de la structure de la surface et des angles de vue et d'éclairage. Une modification de ces caractéristiques peut entraîner des changements spectaculaires dans l'apparence d'un matériau.

Dans cette thèse, nous étudions comment les propriétés optiques peuvent être estimées avec précision à l'aide de l'imagerie informatique. Nous nous concentrons sur les propriétés qui ont un impact sur l'apparence des matériaux. Les solutions actuelles sont basées sur l'utilisation d'instruments qui sont chers, tels que les spectrophotomètres, les profilomètres et les goniophotomètres; par conséquent, nous aimerions explorer l'utilisation d'équipements plus abordables, tels que les caméras. Nous commençons par la lumière réfléchie, et la façon dont la structure et la rugosité de la surface peuvent être mesurées avec du matériel courant tel que des projecteurs et des caméras. Nous continuons avec la lumière transmise, où nous étudions des méthodes de mesure de la transmission hémisphérique, puis nous allons un peu plus loin et proposons une nouvelle conception pour la mesure de la transmission directionnelle.

Des systèmes basés sur la projection de lumière structurée peuvent être utilisés pour une estimation de la profondeur de haute précision. Ces méthodes fournissent des moyens d'estimation de la profondeur sans contact en utilisant une paire de caméra-projecteur calibrée. Le calibrage de ces systèmes est en général une tâche fastidieuse et longue. Nous présentons donc une nouvelle méthode qui permet de calibrer simultanément la géométrie et la radiométrie d'une paire projecteur-caméra. Elle est automatique, simple et ne nécessite pas de matériel spécialisé. Nous pré-déformons et alignons un motif de projection spécialement conçu sur un motif imprimé ayant des propriétés colorimétriques différentes. Après avoir photographié les motifs dans plusieurs orientations, nous effectuons un calibrage géométrique en estimant l'emplacement des coins des deux motifs dans différents canaux de couleur. Nous effectuons un calibrage radiométrique du projecteur en utilisant les informations contenues dans les carrés projetés. Nous montrons que notre méthode est aussi performante que les approches actuelles qui nécessitent toutes un calibrage géométrique et radiométrique séparé. Cela la rend efficace, simple d'utilisation et plus rapide que les techniques de calibrage géométrique et radiométrique spécialisées. Nous appliquons ensuite l'algorithme pour calibrer un système de lumière structurée dans le but d'estimer la structure de surface de surfaces presque planes. La structure de la surface est ensuite utilisée pour calculer la rugosité macroscopique. Sur un ensemble d'échantillons, nos résultats sont similaires à ceux d'un profilomètre commercial de haute qualité.

La mesure précise des propriétés de transmission des matériaux est essentielle dans l'industrie de l'impression et de l'affichage afin de garantir une reproduction précise des couleurs. Nous expérimentons avec différentes géométries de mesure pour la transmission totale et nous montrons que les mesures de transmission sont fortement influencées par la géométrie utilisée, puisque certaines géométries peuvent introduire un biais de mesure. Nous construisons une configuration personnalisée flexible qui peut simuler ces géométries, et l'évaluer à la fois qualitativement et quantitativement sur un ensemble d'échantillons aux propriétés optiques variées. Nous comparons également nos mesures à celles de solutions commerciales largement utilisées et nous montrons que des différences significatives existent entre nos échantillons de test. Cependant, lorsque le biais est correctement compensé, nous n'observons que de très faibles différences. Ces résultats soulignent donc l'importance d'inclure la géométrie de la mesure lors du calcul de la transmission totale.

La fonction de distribution de la diffusion bidirectionnelle (BSDF) est d'une grande importance dans de nombreuses applications, de la génération d'images et des effets spéciaux visuels à la modélisation architecturale de l'éclairage. Nous proposons un nouveau processus pour la mesure de la BSDF, plus précisément la fonction de distribution de la transmission bidirectionnelle (BTDF). Notre approche est basée sur l'imagerie catadioptrique radiale, où nous illuminons l'échantillon mesuré avec un faisceau de lumière collimaté, la lumière émergente de l'autre côté de l'échantillon est ensuite collectée par un diffuseur de lumière cylindrique, et, après de nombreuses réflexions internes, une caméra enregistre la radiance à l'intérieur du cylindre, une fois l'équilibre atteint. Après avoir inversé le transport de la lumière à l'intérieur du cylindre, nous pouvons calculer la BTDF de l'échantillon. Notre approche transforme la complexité mécanique en complexité de calcul. Dans les simulations, le modèle proposée permet d'obtenir des résultats très proches des valeurs correctes, tant pour les BTDF lisses que pour les BTDF à haute fréquence qui présentent des changements d'intensité importants.

Mots clés : imagerie computationnelle, lumière structurée, estimation de la profondeur, transmittance totale, BRDF, BTDF, transport de la lumière, goniophotomètre, imagerie radiale, imagerie catadioptrique.

Contents

Acknowledgements	i
Abstract (English/Français)	iii
List of Figures	xi
List of Tables	xv
Abbreviations and Notation	xvii
1 Introduction	1
1.1 Thesis goals	4
1.1.1 Measuring surface structure of near-planar surfaces	4
1.1.2 Measuring hemispherical and directional transmittance (BTDF)	6
1.2 Thesis outline and main contributions	9
2 Simultaneous geometric and radiometric calibration of a projector-camera pair	13
2.1 Calibration for structured light projection	14
2.2 Related work	15
2.3 Algorithm overview	16
2.4 Design of calibration patterns	18
2.5 Algorithm	19
2.5.1 Image capture	21
2.5.2 Corner detection	21
2.5.3 Prewarp	22
2.5.4 Geometric calibration	22
2.5.5 Radiometric calibration	22
2.6 Experiments	24
2.7 Limitations	27
2.8 Summary	28
3 Surface Roughness Estimation Using Structured Light Projection	29
3.1 Accurate depth estimation with structured light	30
3.2 Structured light patterns for surface profilometry	32
3.2.1 Discrete Coding Methods	32

Contents

3.2.2	Continuous Coding Methods	34
3.3	Experiments	36
3.3.1	Structured light encoding strategy	37
3.3.2	Measurement setup	37
3.3.3	Measurement principle	38
3.3.4	Ground truth measurements	39
3.3.5	Roughness metric	39
3.3.6	Reconstruction errors due to non-linear processing	40
3.3.7	Error compensation	41
3.3.8	Results	42
3.4	Summary	44
4	Comparative analysis of transmittance measurement geometries and apparatus	55
4.1	Related work	55
4.2	Measuring total transmittance	57
4.2.1	Sphere radiance	59
4.3	Our measurement setup	60
4.4	Experiments and results	63
4.4.1	Inter-configuration differences	63
4.4.2	Inter-spectrometer differences	66
4.4.3	Repeatability	67
4.5	Summary	68
5	BTDF Estimation with Radial Imaging	69
5.1	Related Work	73
5.2	Radiometry	76
5.2.1	Geometric quantities	76
5.2.2	Basic quantities	77
5.2.3	The BSDF	79
5.3	The Light Transport Equation	81
5.3.1	Solution to the LTE	82
5.3.2	Surface form of the LTE	84
5.3.3	Discretized surface form of the LTE	86
5.3.4	Analytical solution to the LTE for a simple scene	86
5.3.5	Radiance of a single facet	87
5.3.6	Radiance between all facets	88
5.3.7	Lambertian cylinder	89
5.3.8	Non-Lambertian cylinder	91
5.4	Light transport matrix	92
5.4.1	Computing the light transport matrix	93
5.5	Algorithm	97
5.5.1	Calibration	97
5.5.2	Preprocessing	99

5.5.3	Virtual image	101
5.5.4	Initial facet irradiance from equilibrium radiance	103
5.5.5	BTDF estimation	104
5.6	Simulations	105
5.6.1	Simulations with Mitsuba: A Physically Based Renderer	106
5.6.2	Simulation parameters	107
5.6.3	Single bounce	108
5.6.4	Infinite bounces	112
5.6.5	Compensating for the systematic error	115
5.7	Experiments	117
5.7.1	Experimental setup	117
5.7.2	Image capture	118
5.7.3	Results	119
5.8	Practical considerations and limitations	124
5.9	Summary	125
5.A	Appendix: Shape of \mathbf{F}_r	127
5.B	Appendix: Additional results	128
6	Conclusion	131
6.1	Thesis summary	131
6.2	Future research	133
6.2.1	Simultaneous calibration	133
6.2.2	Surface roughness	133
6.2.3	Parallel goniophotometer	134
	Bibliography	137
	Curriculum Vitae	145

List of Figures

1.1	Material appearance is used for both decorative and functional purposes	2
1.2	Appearance of an object built from materials with different optical properties .	3
1.3	The depth sensor for Apple's Face ID facial recognition system	5
1.4	Systematic reconstruction errors of a flat surface	9
1.5	Surface structure of Canon IJM 618, polyester textile with polymer coating . . .	10
1.6	A measurement setup for transmittance	11
1.7	An isometric view of our proposed design for measuring BTDF	12
2.1	A setup for our simultaneous geometric and radiometric calibration of a projector-camera pair	14
2.2	Systematic reconstruction errors of a flat surface induced by the non-linear processing of the input images by the projector	15
2.3	Flow chart of our proposed method for simultaneous geometric and radiometric calibration of a projector-camera pair	17
2.4	The projected pattern used for geometric and radiometric calibration	18
2.5	Captured image of an aligned superposition of the printed and the projected pattern	19
2.6	The projected pattern in its original size, after having been warped into position atop the printed pattern	23
2.7	Gamma estimation for 14 image calibration set	26
2.8	Systematic errors in the reconstruction of a flat surface	27
3.1	Different instruments for recovering surface structure	30
3.2	An image of a 5 mm × 5 mm section of Canon IJM 618	31
3.3	De Bruijn pattern with $n = 4$ and $k = 3$	33
3.4	10 successive projections that encode the scene with Gray code	34
3.5	Frequency magnitude spectrum of the captured image for FTP	36
3.6	An example printing textile, and the structured-light setup used for the experiments	37
3.7	The measurement principle for our structured-light setup	38
3.8	Source of periodic surface reconstruction errors due to non-linear processing of input signals	40
3.9	Global fabric deformation for the 5 mm × 5 mm Senfa Decoprint Grain scan . .	41

List of Figures

3.10 Surface structure of the flat board used for affixing the samples, measured with our structured-light setup	45
3.11 Surface structure of Senfa Decoprint Grain, polyester textile with acrylic coating	46
3.12 Surface structure of Canon IJM 617, polyester textile with polymer coating . . .	47
3.13 Surface structure of Canon IJM 618, polyester textile with polymer coating . . .	48
3.14 Surface structure of Senfa Decoprint Night, polyester textile with acrylic coating	49
3.15 Surface structure of JM Mediatex Twintex FR, polyester based knitted fabric . .	50
3.16 Surface structure of Senfa Decoprint Grain, and Canon IJM 617, for a 10 mm × 10 mm scan	51
3.17 Surface structure of Canon IJM 618, and Senfa Decoprint Night, for a 10 mm × 10 mm scan	52
3.18 Surface structure of JM Mediatex Twintex FR, for a 10 mm × 10 mm scan	53
4.1 Total transmittance of our diverse set of samples	56
4.2 Comparative methods for measuring total transmittance	58
4.3 Our flexible measurement setup for total transmittance	60
4.4 The set of measurement geometries that we evaluated with our custom setup .	62
4.5 An example set of sample surface texture and appearance	63
4.6 Inter-configuration measurement differences	64
4.7 Inter-spectrometer measurement differences	64
5.1 An example scanning and parallel goniophotometer	70
5.2 Our proposed design for measuring a sample's BTDF	71
5.3 An isometric view of our proposed design for measuring a sample's BTDF . . .	71
5.4 Schematic diagram of the three main working principles of parallel goniophotometers	74
5.5 The Nusselt analog shows the steps of projecting the solid angle visually ¹	77
5.6 Basic radiometric quantities	78
5.7 Schematic of the rendering equation	82
5.8 The three-point form of the LTE	85
5.9 The diffusing cylinder discretized into 5184 equal facets	87
5.10 The outgoing BRDF lobes of Spectralon for incident polar angles of $\theta_i = 0^\circ$ and $\theta_i = 75^\circ$	91
5.11 An iterative computation of a single column of the light transport matrix	96
5.12 Flowchart of our algorithm for BTDF estimation	98
5.13 A captured image with vignetting, a modeled background image, and the captured image after normalizing its pixel intensities with the modeled background image	99
5.14 An image of the sample and cylinder, simulated with Mitsuba	102
5.15 Virtual image and cylinder unrolling	103
5.16 Sample irradiance by a collimated light source, and facet irradiance due to the sample	105

5.17 The physically based renderer Mitsuba allows us to accurately simulate real-world light interactions with objects	106
5.18 The light directions are represented in the spherical coordinate system attached to the sample	107
5.19 Ground truth BTDF of a completely diffusing (Lambertian) sample, and of a Canon IJM 617 sample, a coated white polyester textile	108
5.20 Recovered BTDF after a single light bounce inside the cylinder - smooth	110
5.21 Recovered BTDF after a single light bounce inside the cylinder - high-frequency	111
5.22 Recovered BTDF after an infinite light bounces inside the cylinder - smooth . .	113
5.23 Recovered BTDF after an infinite light bounces inside the cylinder - high-frequency	114
5.24 Recovered BTDF after an infinite light bounces inside the cylinder, and after a compensation has been applied for the systematic error	116
5.25 An implementation of our proposed radial imaging setup	118
5.26 Imperfect alignment of the sample, cylinder and camera, and virtual cylinder projection onto the HDR image	119
5.27 A comparison between the captured HDR images and the images simulated with Mitsuba using the same parameters	120
5.28 Required compensation that will equalize the captured and the simulated images	121
5.29 The BTDF ratio when computed from the captured images, to the ground truth	123
5.30 Practical limitations of our proposed setup for estimating BTDF	125
5.31 Recovered BTDF after an infinite light bounces inside the cylinder, and after a compensation for the systematic error - Canon IJM 601 and IJM 637	128
5.32 Recovered BTDF after an infinite light bounces inside the cylinder, and after a compensation for the systematic error - Avery Dennison MPI 2020 and MPI 3151	129
6.1 Geometry for capturing a different outgoing hemisphere	134
6.2 Geometry for capturing an oblique incoming direction	135
6.3 Schema of the mirrored cylinder	136
6.4 Simulated images for a diffuse sample and mirrored cylinders	136

List of Tables

2.1	Mean reprojection errors (RMSE in pixels) for camera and projector calibration and the number of images per orientation as reported by the authors	25
3.1	The characteristics of the textile samples that we used for testing	36
3.2	A summary of the roughness results	43
4.1	The measurement configurations that we used to capture total transmittance with our setup	61
4.2	Inter-configuration measurement differences	65
4.3	Inter-spectrometer measurement differences	66
4.4	Updated inter-spectrometer measurement differences	67
4.5	Measurement repeatability	67
5.1	Summary of the measurement configurations that we evaluated in this work . .	72
5.2	Summary of the parameters used to generate the virtual image in the simulations and in the experiments	101
5.3	Experimental setup, camera, rendering and virtual image parameters for the Mitsuba simulations and for the experiments	109
5.4	Summary of the camera parameters used for capturing photos for the experiments	118

Abbreviations and Notation

List of Abbreviations

Abbreviation	Description
BRDF	Bidirectional Reflectance Distribution Function
BSDF	Bidirectional Scattering Distribution Function
BTDF	Bidirectional Transmittance Distribution Function
CMOS	Complementary Metal-Oxide-Semiconductor
CSL	Coded Structured Light
DLP	Digital Light Processing
DRA	Diffuse Reflectance Accessory
FTP	Fourier Transform Profilometry
FWHM	Full Width at Half Maximum
GPU	Graphics Processing Unit
HDR	High-Dynamic Range (image)
IS	Image Stabilization
LCD	Liquid Crystal Display
LED	Light Emitting Diode
LTE	Light Transport Equation
LTM	Light Transport Matrix
MAE	Mean Absolute Error
ND	Neutral-Density (filter)
PCA	Principal Component Analysis
RMSE	Root Mean Square Error
SL	Structured Light
SNR	Signal-to-Noise Ratio
USM	UltraSonic Motor
UV	UltraViolet

List of Symbols

Symbol	Description	Page
i	<i>index</i>	general
i, o	description (e.g., incident, outgoing)	general
p	point	general
p, F, L	<i>quantity</i>	general
\mathbf{x}	3D homogeneous coordinate vector	20
\mathbf{X}	4D homogeneous coordinate vector	20
\mathbf{R}	rotation matrix	20
\mathbf{t}	translation vector	20
\mathbf{K}	intrinsic camera parameters	20
\mathbf{H}	homography	20
γ	non-linear tone mapping function	23
\mathbb{E}	expected value	23
S_q	root mean square height of a surface	39
S_z	maximum height of a surface	43
L	radiance	59
Φ	radiant flux (power)	59
M	sphere multiplier	59
ρ	surface reflectance	59
f_i	fraction of port i to the total sphere area	59

Symbol	Description	Page
p	point in 3D space	76
ω	normalized direction on the unit sphere	76
\mathcal{S}^2	sphere of directions	76
$d\omega$	solid angle, or a set of directions on the unit hemisphere	76
$d\omega^\perp$	projected solid angle	76
\mathbf{n}_p	surface normal at point p	76
$E(p)$	irradiance of a surface at point p	78
$M(p)$	radiant exitance from a surface at point p	78
$L(p, \omega)$	radiance arriving at/leaving from point p from/in direction ω	79
$f_s(p, \omega_i, \omega_o)$	BSDF (Bidirectional Scattering Distribution Function)	79
$f_r(p, \omega_i, \omega_o)$	BRDF (Bidirectional Reflectance Distribution Function)	80
$f_t(p, \omega_i, \omega_o)$	BTDF (Bidirectional Transmittance Distribution Function)	80
\mathcal{H}^2	hemisphere of directions about a surface normal	80
$\mathbf{x}_A(p, \omega)$	ray-casting function from point p in the direction ω	81
\mathcal{G}	propagation operator	82
\mathcal{K}	local scattering operator	83
\mathcal{T}	light transport operator	83
\mathcal{S}	solution operator	83
\mathcal{I}	identity operator	83
$L(p' \rightarrow p)$	outgoing radiance from point p' towards point p	85
$f_s(p'' \rightarrow p' \rightarrow p)$	surface form of the BSDF	85
$G(p \leftrightarrow p')$	geometric coupling between points p and p'	85
$V(p \leftrightarrow p')$	visibility between points p and p'	85
$\mathbf{l}^{\rightarrow p_j}$	outgoing radiance of all points towards p_j	87
$\mathbf{f}_r^{p_j \rightarrow p_i}$	BRDF for light incident from all p , reflecting from p_j , towards p_i	87
$\mathbf{G}^{\rightarrow p_j}$	projected solid angle of all facets at point p_j , diagonal matrix	87
$\mathbf{g}^{\rightarrow p_j}$	projected solid angle of all facets at point p_j , vector	87
$\mathbf{l}^{p_j \rightarrow}$	outgoing radiance from p_j towards all points	88
$\mathbf{l}_e^{p_j \rightarrow}$	radiance emitted from p_j towards all points	88
$\mathbf{F}_r^{p_j}$	BRDF: light incident from all p , reflecting from p_j , towards all p	88
\mathbf{l}	outgoing radiance from all points towards all points	88
\mathbf{l}_e	emitted radiance from all points towards all points	88
\mathbf{F}_r	BRDF: light incident from all p , reflecting from all p , towards all p	88
\mathbf{G}	projected solid angle of all facets at all points p , matrix	88
\mathbf{g}	projected solid angle of all facets at all points p , vector	88
\mathbf{I}	identity matrix	89
\mathbf{G}^L	projected solid angle of all facets at all points p , compact matrix	90
\mathbf{e}_0	initial irradiance of all facets by the measured sample	93
\mathbf{T}	light transport matrix	94

1 Introduction

How we visually perceive non-emissive objects in our surroundings depends on the interaction of light with the materials that comprise them according to their optical characteristics. During the interaction, the light is reflected, transmitted, or absorbed by the materials. These optical properties are determined by the internal composition of scattering and absorbing components, the microscopic surface structure, and the materials' refractive indices.

Material appearance is more subjective, and it is related to our visual perception. We can observe that an object is a given color, whether it is matte or glossy; opaque, translucent or transparent, etc. The material appearance depends on the optical properties of the materials, but also on the illuminating and viewing directions. In addition, macroscopic surface structures can influence the appearance through shadowing and interreflections.

Sight is arguably one of our most developed senses. It makes us capable of accurately identifying the material composition, characteristics, and function of objects based solely on their appearance [25]. We can, for example, identify if something is made out of paper, plastic, textile, wood, metal, and within those broad classes we can further distinguish different materials, e.g., cotton, wool, or silk. Furthermore, we can tell if an object is matte, translucent, solid, sticky, slippery, expensive, or if it is edible, made to be worn, or if it is a tool. We are very good at learning visual associations. Based on the appearance subtleties, we can generally tell if a photograph was taken 10 years ago, in the 1990s, or in the 1960s, irrelevant of the scene content.

The reproduction and simulation of material appearance is an active area of research in the scientific community, but it is also very relevant to industry. Its importance naturally leads to companies investing significant resources in mastering the look of products. To stand out, their products should look, but also feel, high quality, or even expensive. The International Commission on Illumination (CIE) describes its importance as: "Visual appearance can be one of the most critical parameters affecting customer choice and it needs, therefore, to be quantifiable to ensure uniformity and reproducibility."¹

¹<http://cie.co.at/publications/framework-measurement-visual-appearance>



(a) Car paint



(b) Felt hat



(c) Wristwatch



(d) Morpho menelaus butterfly

Figure 1.1 – Material appearance is used for both decorative and functional purposes. (a) Car paint with pearlescence. (b) Felt is isotropic, and diffuses light well. (c) Wristwatch made out of materials with different optical properties, and surface finishes. (d) Butterfly with iridescent wings. Sources: (a) Saud Al-Olayan, (b) Paula Cooper, (c) Bruno van der Kraan, (d) Sandra.²

Example products for which material appearance is especially important, are shown in Figure 1.1. The paint on the car from Figure 1.1a is pearlescent, and its color gradually changes as the viewing and illumination directions change. Material appearance is also very important in the clothing industry, where designers mix and match materials to achieve a certain look. Similarly, the watch industry is constantly experimenting with different materials, and surface finishes to make their products appear more luxurious. The final Figure 1.1d shows the Morpho menelaus butterfly. It is very interesting, because the way it achieves its iridescent blue color is not with pigments, but with nanometre scale gratings. Only the blue light is reflected, while the other wavelengths of the incident light are canceled out by interference.

²<https://www.flickr.com/people/69929929@N06/>, <https://www.flickr.com/people/peacooper/>,
<https://unsplash.com/@brunovdkraan>, <https://www.flickr.com/people/supervliegus/>




Figure 1.2 – Appearance of an object when built from materials with different optical properties. In all images, the object is observed from the same direction. The visual differences between the images are caused by differences in the internal composition of absorbing and scattering components, surface structure, and in the illumination directions. (a)-(d) are illuminated from the same directions (the environment map can be best seen in (c)), whereas (e) and (f) are illuminated from different directions. (c)-(f) are anisotropic samples. Images taken from: <http://rgl.epfl.ch/materials>.

In Figure 1.2 we can see an example object simulated as if it was made from different materials. The images were generated with a rendering software that respects physically accurate interactions between light and matter. The materials were measured from real samples with different optical characteristics, which results in striking appearance differences between the images. The material in Figure 1.2a is a white plastic, and it is the most diffusely reflecting material available.³ The acrylic felt in Figure 1.2b is also isotropic, however, its internal absorbing components give it a purple color. The aluminium object from Figure 1.2c exhibits fairly precise reflections due to its smooth surface structure. The three images in the bottom row (Figure 1.2d-1.2f) represent the same material, the wing of a butterfly (Morpho menelaus), under different illumination conditions. It is an anisotropic material, and therefore the appearance of the object changes significantly when the illumination arrives from different directions.

³ <https://www.labsphere.com/labsphere-products-solutions/materials-coatings-2/targets-standards/diffuse-reflectance-standards/diffuse-reflectance-standards/>

Fast, accurate and reliable measurements of the optical properties and the material appearance are key. They enable faster product development cycles, maximize the appearance quality by allowing more iterations within a design cycle, but are also important for quality control purposes. Over time, we have developed numerous techniques, devices, and industrial standards for the purpose of measuring the optical properties of materials. For example, to measure reflectance or transmittance, we use spectrophotometers. If we are interested in measuring the angular distribution of light, we use goniophotometers. Profilometers are used for measuring the surface structure and roughness of materials. These devices are more or less accurate and reliable, but they are generally expensive, and therefore accessible to a select group of institutions. To illustrate this, some of the largest publicly available BRDF datasets count less than 200 samples.⁴

1.1 Thesis goals

In this thesis, we explore alternative techniques for measuring the optical properties of materials. Our methods are based on computational imaging, and revolve around cameras and projectors, as these are generally more affordable devices. In the first part of the thesis, we look into measuring the properties that influence the reflected light. Namely, we explore the accuracy of a well calibrated projector and camera pair for measuring the surface structure and roughness of near-planar surfaces (Chapters 2 and 3). In the second part, we turn our attention to measuring transmittance. We first explore how to accurately measure total transmittance (Chapter 4), and then evaluate a novel design for the measurement of the Bidirectional Transmittance Distribution Function (Chapter 5).

1.1.1 Measuring surface structure of near-planar surfaces

The advances in virtual reality display technology and the emergence of 2.5D/3D printers renewed the research interest in 3D acquisition of real world objects and scenes. Digitization for virtual reality applications is often concentrated on capturing our immediate surroundings, on the scale from simple objects to complete buildings. These applications put the primary focus on acquisition speed and robustness to different types of illumination and scales, where the measurement precision is generally of secondary importance. The interest that virtual reality applications created can be seen through the large number of commercially available products and prototypes, some of which include Microsoft Kinect, Intel RealSense, Occipital Structure Sensor, Asus Xtion, and HP Pro S3. Modern smartphones often have face unlocking features, that scan the face with structured light projection in the invisible near-infrared radiation. Figure 1.3 shows the depth sensor for Apple's Face ID technology, which is based on structured light projection.

⁴<https://www.merl.com/brdf/>, <http://btf.utia.cas.cz>, <http://rgl.epfl.ch/materials>



Figure 1.3 – The depth sensor for Apple’s Face ID facial recognition system. It works by projecting infrared dots onto a face, which are captured and decoded to produce a 3D map. Image source: Apple.

Structured light systems are popular for indoor 3D scanning of objects, since they are relatively easy to build, and there are scene encoding methods for almost any scene complexity. They are built with a minimum of a single camera and a single projector. Passive stereo systems are very similar in nature, the main difference being that they are built using two cameras. The role of the projector is to encode the scene with known light patterns, such that reliable correspondences can be established between the camera and projector pixels. A triangulation procedure then projects rays from the corresponding pixels in the two devices, and determines their intersections in 3D space. For indoor scenes that often lack texture, this is an advantage over passive stereo setups, built from two or more cameras.

Geometric calibration is an essential part of methods for 3D scanning that rely on triangulation. Given a set of cameras and projectors, it determines their internal and external parameters. The internal parameters are used for projecting rays that start at the optical center of the device, cross through a pixel of interest in the image, and continue in the space in front of the device. The external parameters determine the relative positions and rotations between the set of devices. Clearly, it needs to be performed each time there is a change in the geometry of the setup.

Radiometric calibration characterizes the output intensities of a devices as a function of the input intensities. Non-linear processing is almost always present in the imaging pipelines, since the devices are made to comply with given color spaces (e.g., sRGB), which in turn optimize for the non-linearity of perceived stimuli by the human visual system. For common cameras, this non-linear processing can often be disabled, yielding linear images. However, it is often impossible to disable it for displays and projectors.

To facilitate the use of inexpensive projectors for depth estimation with structured light, in Chapter 2 we propose a method that combines the geometric and the radiometric calibration of a projector-camera pair. Our method optimizes and then leverages the colorimetric properties of the calibration charts to separate the different kinds of information needed for the two calibrations, which speeds-up the calibration process. The method can be used off-the-shelf by the industry, and could be easily extended to cameras with more than 3-channels (e.g., multi-spectral cameras).

The surface structure is of great interest for our desire to master the appearance of products. If we take printing as an example, very accurate color prediction models for printing on flat diffusing substrates have been developed over the past decades. However, as digital printing on textiles becomes more popular due to faster production and cost-effective print runs, these models fail to provide accurate color management in the presence of the macroscopic surface roughness of the textile. To overcome this obstacle to a wider adoption of digital printers, new color prediction models will have to be devised, which take into account the local interreflection of light, and provide predictions with directional dependence.

Therefore, in Chapter 3, we focus on precise and accurate depth estimation with a structured light scanner built from mainstream devices. Our goal is to estimate the surface roughness of near-planar surfaces, such as printing textiles. We provide a transparent analysis of the problems common to this reconstruction method, and also give advice regarding how the results could be improved. We believe that this is very valuable, as it will help other researchers and practitioners build more robust acquisition setups by minimizing the errors inherent from the reconstruction method.

1.1.2 Measuring hemispherical and directional transmittance (BTDF)

We continue to transmitted light, where we study methods for measuring hemispherical transmittance, and then go one step further and propose a novel design for the measurement of directional transmittance. This is an active area of research that still needs to be properly standardized, as the CIE writes: “Translucency is a subjective term that relates to a scale of values going from total opacity to total transparency. This whole subject area needs investigation to find a rigorous measurement solution that will probably be industry specific.”⁵

Transmittance measurement is used in a wide range of industries, from solar cells and light filter manufacturing to quality control in food production. The accurate characterization of transmittance is especially important for applications that depend on backlight illumination, which are created by depositing an image onto a substrate that is illuminated from the back side. Common backlit devices include LCD displays and printed advertisements. The correct display on these devices requires accurate spectral prediction models, and transmittance measurements are used extensively in their calibration.

⁵<http://cie.co.at/publications/framework-measurement-visual-appearance>

In Chapter 4 we present different measurement geometries for total (hemispherical) transmittance. We explain where the measurement differences stem from, and compare the results meticulously. We believe that this chapter is of great utility not only to the scientific, but also to the industrial community, as it demonstrates how important it is to fully disclose the type of measurement instrument used. To highlight this point, the CIE recently created a Technical Committee which focuses solely on transmittance measurement:⁶

“This TC will investigate how to improve the accuracy of the conventional haze measurement methods and develop new techniques to obtain total transmittance, diffuse transmittance, and transmittance haze in a single sequence of measurements. In addition to addressing the problem of measurement inconsistency caused by employing different measurement methods or different measurement instruments which has perplexed industry for some time, the study of high haze (> 40 %) measurement is also included. The purpose of this TC is to give guidance to industry for best practice metrology, to enable them to develop and manufacture consistent and reliable products.”

Computer generated simulations and animations require even more granular data to achieve visual realism. A higher resolution representation of the transmittance includes a measurement for every incoming and outgoing direction on the hemispheres of directions, which, for transmittance, lie on the opposite sides of the measured sample. This is also known as directional transmittance, and the hemispherical transmittance from above can be computed by integrating over the hemisphere of outgoing directions for a single incoming direction. The mathematical description of how light is scattered after an interaction with a surface is more commonly known as the Bidirectional Reflectance or Transmittance Distribution Function (BRDF and BTDF). The proper measurement of these functions is a very active area of research. In May 2019, a joint research project, titled “New quantities for the measurement of appearance”, was launched as a collaboration between National Metrology Institutes, partners from Industry and the research Community in Europe.⁷ One of the main goals of the project is to define the new quantity BTDF, i.e., to standardize its mathematical definition.

Physically-based rendering relies heavily on accurate BRDF and BTDF data for generating realistic ray-traced images. More recently, computer hardware became powerful enough to perform real-time ray tracing. Major graphics cards manufacturers are now including ray-tracing capabilities to their newest models of graphics cards. These capabilities will also be present in the very popular gaming consoles, which makes it a mainstream technology that is here to stay. As the hardware capabilities evolve, digital artists and game designers will require BTDF measurements of increasingly more intricate materials to captivate their audiences. Cultural heritage preservation, visualization and reproduction is another area where this

⁶<http://cie.co.at/technicalcommittees/measurement-total-transmittance-diffuse-transmittance-and-transmittance-haze>

⁷<https://bxdiff.cmi.cz>

Chapter 1. Introduction

data is of great importance. It also has more functional applications, such as architectural simulations of indoor illumination.

In Chapter 5, we present a novel design for a parallel goniophotometer for measuring the BTDF. Our approach is based on computational imaging, where we invert the radiometric model to compute the radiance that leaves the measured sample. It is a unique method which trades mechanical complexity with computational complexity, where prior methods were mainly optimizing the mechanical designs of the devices. Our design is based on conventional and affordable optical elements, and can be scaled up or down, depending on the application.

1.2 Thesis outline and main contributions

In this section we present a condensed version of the thesis chapters, and summarize their main contributions.

Chapter 2: Simultaneous geometric and radiometric calibration of a projector-camera pair

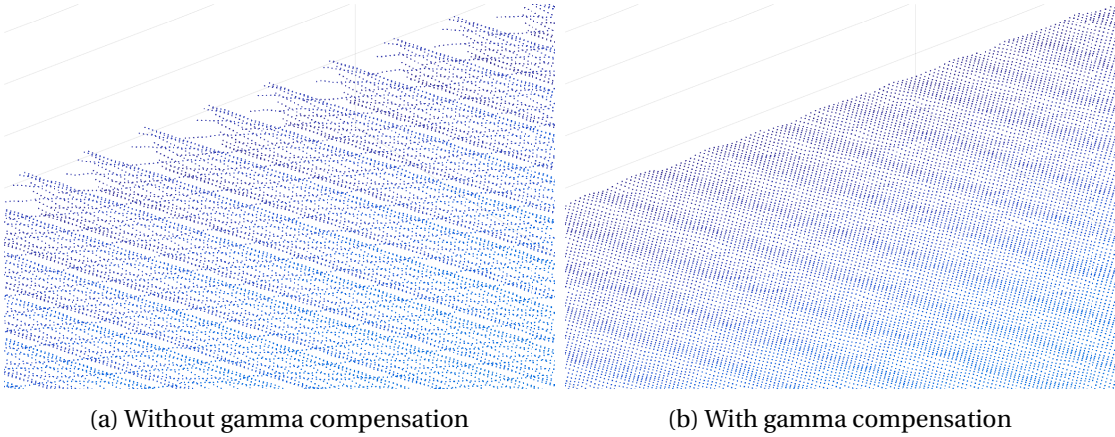


Figure 1.4 – Systematic reconstruction errors of a flat surface induced by the non-linear processing of the input images by the projector.

In this chapter we propose a novel algorithm for geometric calibration of a projector-camera pair that also provides a radiometric calibration of the projector. The algorithm works similarly to the existing algorithms for geometric calibration, in the sense that it does not require additional hardware or user intervention. The added value that it offers is due to the special design of the calibration patterns, and simple, yet effective online computation. We prewarp and align a specially designed projection pattern onto a printed pattern of different colorimetric properties. After capturing the patterns in several orientations, we perform geometric calibration by estimating the corner locations of the two patterns in different color channels. We perform radiometric calibration of the projector by using the information contained inside the projected squares. We show that our method performs on par with current approaches that all require separate geometric and radiometric calibration. This makes it efficient, user friendly, and faster compared to dedicated geometric and radiometric calibration techniques.

Summary of contributions in Chapter 2:

- We propose a novel algorithm for geometric calibration of a projector-camera pair that also provides a radiometric calibration of the projector.
- We show that the proposed algorithm provides high quality calibration by comparing it to current state-of-the-art methods, while being more time efficient compared to dedicated geometric and radiometric calibration techniques.

Chapter 3: Surface roughness estimation using structured light projection

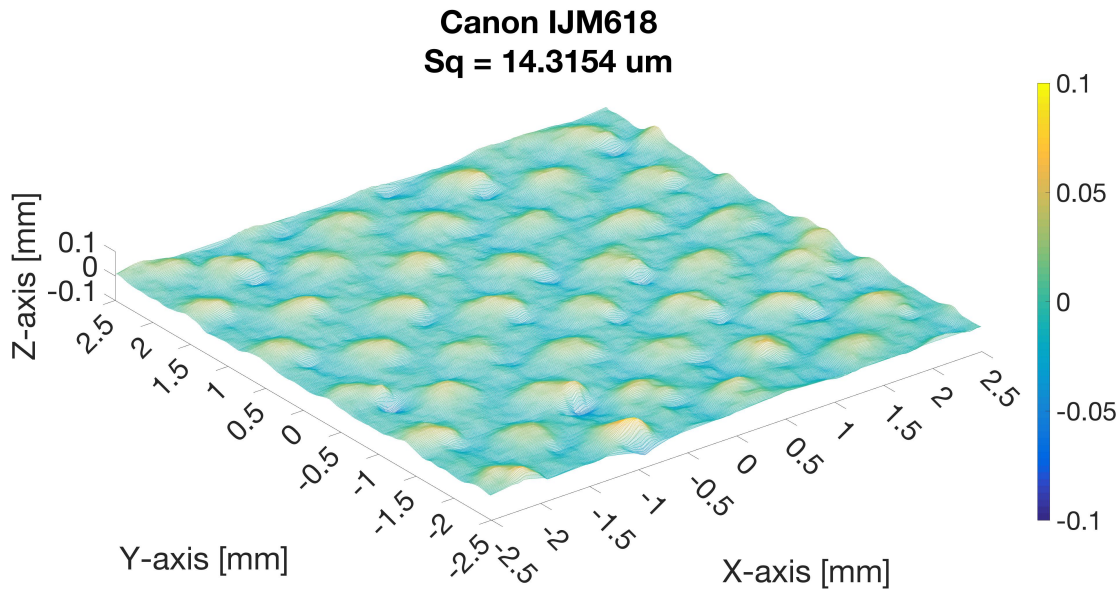


Figure 1.5 – Surface structure of Canon IJM 618, polyester textile with polymer coating. The colormap on the right side color codes the height of the surface points. The height is represented as the deviation from the best fitting plane to the scanned point cloud.

In this work we leverage the algorithm presented in Chapter 2. We begin with an overview of the main categories of structured-light (SL) encoding techniques, and discuss the potential sources of errors. We build a SL system and calibrate it with the algorithm from Chapter 2. We use the structured light setup to estimate the surface structure of printing textiles. Our scans are compared to those from a commercial profilometer based on the confocal principle, both qualitatively and quantitatively by computing the surface roughness measure. When properly calibrated, we achieve comparable results to the profilometer on samples with moderately complex surfaces. The SL setup offers faster scanning and/or wider scanning area than the profilometer, at the expense of lower resolution.

Summary of contributions in Chapter 3:

- We discuss the main categories of structured light encoding strategies, and discuss their potential sources of errors.
- We show that a properly calibrated structured light system can provide comparable surface roughness estimates to a profilometer over a set of measured samples.

Chapter 4: Comparative analysis of transmittance measurement geometries and apparatus




Figure 1.6 – A measurement setup for transmittance. It consists of an illuminating optical fiber, an iris diaphragm, an integrating sphere and a capturing optical fiber connected to a spectrometer. Depicted is the configuration $0^\circ:d$.

In this chapter, we turn our attention to the transmitting properties of materials. In comparison with reflectance measurement, where the impact of different geometries ($0^\circ:45^\circ$, $d:8^\circ$) has been thoroughly investigated, there are few published articles related to transmittance measurement. We therefore summarize different techniques for measuring total transmittance, and show the sources of the differences. We also build a system for total transmittance measurement to evaluate several of the measurement geometries, and compare our measurements to those from commercial solutions. Certain measurement geometries are found to introduce a significant bias, which should be compensated. Those findings stress the importance of including the measurement geometry when reporting total transmittance.

Summary of contributions in Chapter 4:

- We investigate different measurement geometries for total transmittance.
- We show that the transmittance measurements are highly affected by the geometry used, since certain geometries can introduce a measurement bias.
- We compare our measurements against those of widely used commercial solutions, and show that significant differences exist over our test set.

Chapter 5: BTDF estimation with radial imaging




Figure 1.7 – An isometric view of our proposed design for measuring a sample’s BTDF. The elements have been cut in half for clarity.

Going further with transmittance measurement, we introduce a novel design for a parallel goniophotometer. It is based on radial catadioptric imaging. The design is very simple, and it trades mechanical complexity with computation. Namely, it consists of a diffusing cylinder with the measured sample at one of its openings, and the camera at the other. The sample is illuminated by a collimated beam of light, the cylinder collects the emerging light, and after many interreflections from the cylinder walls, the camera captures the steady-state radiance of the cylinder interior. We then reverse the light transport inside the cylinder to compute the initial radiance that leaves the sample. From there, we can compute the directional distribution of light, i.e., the BTDF.

We present a rigorous discussion on the mathematical foundation of the method. We then analyze it through simulations inside Mitsuba, a physically based renderer. The simulations show a good alignment of the computed BTDF and the ground truth, for smooth functions, but also for functions with abrupt changes. We also build an experimental setup out of off-the-shelf optical components. However, due to misalignment issues of the optical equipment, the results are not entirely satisfactory. We believe that an industrial grade implementation could provide results close to the ones in the simulation.

Summary of contributions in Chapter 5:

- We propose a novel method for the measurement of BTDF.
- We provide a rigorous mathematical discussion on the foundation of the method.
- We analyze the method through simulations, and show that our design provides accurate results over a range of BTDFs.

2 Simultaneous geometric and radiometric calibration of a projector-camera pair

3D sensing is gaining momentum with the increase of computational power and the possibility to display and fabricate the results with high fidelity. Structured light (SL) systems are among the most commonly used for 3D object scanning because they can easily be built using off-the-shelf components. In this chapter, we present a novel method for joint geometric and radiometric calibration of a projector-camera pair, which form an SL setup, that is also known as an active stereo setup. Our method is simple, does not require specialized hardware, and calibration is achieved in less than one minute.

In Section 2.1, we discuss the role of calibration in an SL setup for depth estimation. While geometric calibration is essential for computing the point cloud of the scanned scene, we also identify the important case where radiometric calibration is necessary. Namely, radiometric calibration should be performed when the camera and/or projector apply non-linear processing, and the selected structured light encoding methods feature codes with continuous intensity variation. It is then followed by a discussion of the related work in Section 2.2.

Section 2.3 provides an overview, and Section 2.5 presents the details of our simple, yet accurate algorithm for calibration. We prewarp and align a specially designed projection pattern onto a printed pattern of different colorimetric properties. After capturing the patterns in several orientations, we perform geometric calibration by estimating the corner locations of the two patterns in different color channels. We perform radiometric calibration of the projector by using the information contained inside the projected squares. The design of the printed and projected patterns is presented in Section 2.4.

The experimental part then follows in Section 2.6, where we test both the geometric calibration of the camera-projector pair, and the radiometric calibration of the projector. We show that our method performs on par with current state-of-the-art methods that all require separate geometric and radiometric calibration, while being more efficient and user friendly. We conclude with a discussion of the proposed method's limitations and a summary.¹

¹The material in this chapter was presented at CVPR 2017 [90].

2.1 Calibration for structured light projection

A basic SL setup is made up of a single camera and a single projector that are *geometrically* calibrated (see Figure 2.1). The depth estimation principle is analogous to a dual camera stereo system, where one camera is replaced by a projector. This is possible since the geometric image formation of the projector is equivalent to that of the camera. The role of the projector is to encode the measured scene with a known pattern that can be easily decoded and identified in the camera image, allowing for dense pixel correspondences.

A thorough classification and explanation of the structured light coding techniques is presented by Salvi et al. [87]. Discrete coding methods are generally constructed through the projection of a limited set of highly contrasting intensities. If we take binary time-multiplexing codes as an example, the value that each pixel can take is either 0 or 1, as these are the most contrasting values. On the other hand, the continuous coding methods can use the full range of intensity values (e.g., sinusoidal waves), where each pixel can take a value in the range $[0, 1]$. These methods require the projection and capture of linear intensity signals, and any non-linear shift between the input values and the decoded values leads to shifted correspondences between the projected and captured images, illustrated in Figure 2.2a. Therefore, in addition to geometric calibration, the continuous coding methods also require a *radiometric* calibration.

Radiometric calibration models the input to output intensity response of a device, because its true values are seldom known beforehand. While linear intensity images are in general easy to obtain with today's cameras, the projected images normally undergo non-linear processing imposed by the projector. The most common non-linear processing consists of applying a global tone-mapping power-law expression (also called gamma expansion) to the input intensities before projection. To linearize the projected images, the input images need to be subjected to gamma compression.




Figure 2.1 – A setup for our simultaneous geometric and radiometric calibration of a projector-camera pair. The underlying printed pattern is shown in the inset.




Figure 2.2 – Systematic reconstruction errors of a flat surface induced by the non-linear processing of the input images by the projector. For this example, a conventional phase-shifting encoding was used [87].

Geometric calibration of a projector-camera pair has to be performed on a regular basis, especially while changing the parameters of the system (baseline, focal length, etc.) or if the measurement setup needs to change locations. Also, for setups that are built with lower grade materials and tools, calibration needs to be performed more often. Radiometric calibration can become invalidated through time, depending on the current temperature of the projector and age of the lamp. Therefore, we are strongly incentivized to design simple and fast calibration procedures that are easy to use even outside of research environments.

2.2 Related work

Geometric calibration The research on geometric calibration of projector-camera pairs has produced several classes of methods: two step methods [7, 35], juxtaposed methods [8, 15, 46], methods based on structured light projection [57, 68, 99, 101, 105], and methods that leverage different color channels [7, 105]. All projector-camera approaches require a printed chart that carries features with known 3D locations, and a projected chart or a projector-camera pixel correspondence map that is used for the calibration of the projector.

Two step methods rely on first calibrating the camera and then using the calibrated camera to calibrate the projector. These approaches are obviously time consuming since two sets of images have to be acquired. The *juxtaposed methods* position the printed and the projected patterns such that the interference of the markers of the two patterns is minimized. Correct positioning is achieved by either prewarping the projected pattern by a homography that positions it in the correct location when projected, or by projecting it right next to the printed pattern such that there is no overlap. When positioning the two patterns next to each other, a wider field of view camera lens is required, which leaves certain regions of the camera image not utilized when performing the depth estimation by SL projection. The methods that utilize

Chapter 2. Simultaneous geometric and radiometric calibration of a projector-camera pair

prewarping require the calibration board to be relatively static for at least two consecutive image acquisitions as the prewarp is invalidated by movement.

The methods based on *structured light* projection tend to establish dense projector-camera correspondences, i.e., they try to relate each projector pixel to a camera pixel. The SL is projected over a printed calibration chart, whose features are later used for the calibration of both the camera and the projector. Robust SL codes require the projection of a multitude of patterns, and this process has to be repeated for each orientation of the printed calibration board. This makes them slow, data intensive, and require an immobilization of the calibration chart for the duration of the SL encoding. The methods that leverage different *color channels* design printed and projected calibration charts of complementary colors that minimize the crosstalk between the two charts. One obvious drawback of this class of methods is the color channel crosstalk that can lower the contrast and therefore the precision of the corner detection. In this work, we provide a better colorimetric encoding strategy.

Radiometric calibration Prior work on radiometric calibration mostly focuses on cameras. Debevec et al. [18] describe an approach that estimates the response curve of a camera by taking multiple images of a static scene with different amounts of exposure. Subsequent works [58, 67], propose different models for the response curve (gamma, polynomial, non-parametric, PCA based). These works are not directly applicable to a projector calibration, as projectors cannot measure scene radiance, however they do provide valuable insight into the complexity and effectiveness of different techniques. In [72], Nayar et al. perform radiometric calibration of a projector by projecting a succession of different intensity images. They later store the camera responses to those inputs in an inverse lookup table. We adopt a similar approach, however, we estimate a gamma function from the simultaneous display of different intensities.

Given the above classification of methods for geometric calibration, our approach can be classified as a juxtaposed method that leverages different color channels. This fusion, together with an innovative projection pattern design (discussed in Section 2.4), facilitates the geometric calibration and allows for simultaneous radiometric calibration.

2.3 Algorithm overview

An overview of our algorithm is provided in Figure 2.3. For each orientation of the calibration board we consecutively execute steps 1 and 2. We start by projecting a specially designed checkerboard pattern, described in detail in Section 2.4, on top of a conventional checkerboard pattern that is printed with a complementary color to the projected one. In step 1, we capture an image of the arbitrarily aligned patterns and detect their corners. This allows us to compute a composition homography that relates the projector image to the printed pattern. We use this homography to compute a prewarp of the projected pattern, such that when projected it lands right on top of the printed pattern, with a shift of half a printed square. Prewarp has

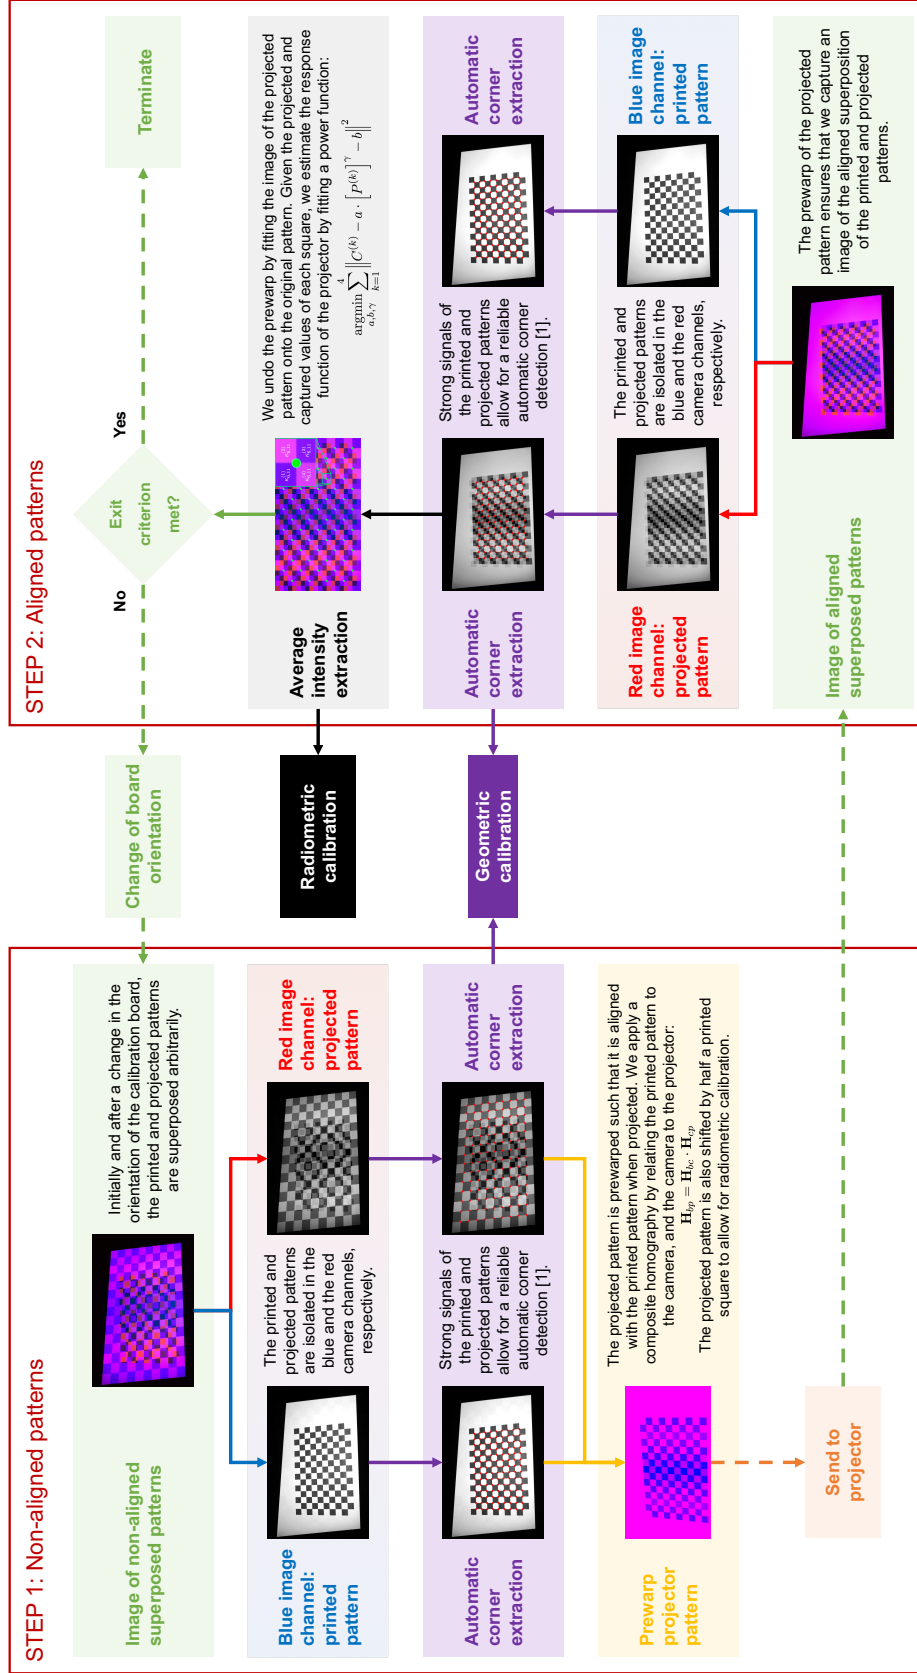


Figure 2.3 – Flow chart of our proposed method for simultaneous geometric and radiometric calibration of a projector-camera pair. Each orientation of the calibration board requires two images. In step 1 we start by capturing an image of the arbitrarily aligned printed and projected patterns. We detect the corners of the two patterns and compute a prewarped that aligns them. In step 2, we capture an image of the aligned patterns. This allows us to extract projected intensities for radiometric calibration in addition to the corner information used for geometric calibration.




Figure 2.4 – The projected pattern used for geometric and radiometric calibration. The red pixel intensities (in %) of the squares vary across the chart, which is leveraged for the radiometric calibration of the projector. The values in the top-right corner are example intensities in the red channel. The blue pixel intensities are constant.

already been employed by Audet et al. [8] and Chen et al. [15] for camera-projector geometric calibration, however, none of those methods perform radiometric calibration.

At step 2, we capture an image of the aligned patterns. After the printed and projected patterns are aligned and shifted by half a square, we extract the information encoded in the projected squares that land on white (empty) squares of the printed pattern. This information encodes the projector response to a range of different input values, which we then use for radiometric calibration.

The above process is repeated for different calibration board orientations until the desired number of loops has been acquired, usually between 10 and 15.

2.4 Design of calibration patterns

Printed pattern The printed pattern is a standard flat checkerboard pattern comprised of white and yellow squares. Its configuration is equivalent to that of the projected pattern (Figure 2.4), but it features only fulltone squares. An example printed pattern can be seen in the inset of Figure 2.1. We opted to use yellow squares because the yellow colorant is almost transparent to the red and green wavelengths, but it absorbs in the blue wavelengths. Therefore, the yellow squares show as dark squares in the blue channel of the camera, where we can easily detect the checkerboard pattern (see Figure 2.5b). While [7] have used cyan and white squares, we opted for a yellow-white configuration, because the real yellow inks are in general much closer to the perfect yellow ink than cyan inks are. Yellow ink is much more transparent in the green-red wavelength range (500-700 nm) than the cyan is in the blue-green range (400-600 nm). This translates into a fairly visible ghost of the cyan pattern in the blue image channel - the channel where the projected pattern should be isolated.




Figure 2.5 – Captured image of an aligned superposition of the printed and the projected pattern. This is the input to step 2 of the algorithm. Corner detection for the two patterns is performed in different color channels.

Projected pattern The projected pattern is shown in Figure 2.4. It is specifically designed such that we can simultaneously perform both geometric and radiometric calibration of the projector. Note that the squares are colored with a blue-magenta palette. More precisely, all squares are composed of 100 % blue and varying intensities of red. The top-right corner of Figure 2.4 shows sample intensity values for the red color channel. The patches that have higher red intensity become increasingly magenta due to the additive mix with blue. The pattern is symmetric across the middle column, and the intensity values alternate across rows.

These colors are chosen for two reasons. First, the high intensity blue is present to illuminate the printed pattern. Since the yellow squares absorb in the blue region, this increases the contrast of the printed pattern and makes the calibration independent of external illumination. Second, the squares include a red component so that we can detect the projected checkerboard corners in the red channel of the captured images where the printed pattern is rather transparent (Figure 2.5c). Although the squares have varying intensities of red, with our design we preserve the contrast ratio of each corner; we impose an intensity difference of at least 40 % between the dark and the bright squares at a corner.

The varying intensities of red across the projected squares are used for the radiometric calibration of the projector. In the current design we have 11 intensity gradations, ranging from 0 % to 100 %, in 10 % steps.

2.5 Algorithm

Geometric calibration establishes a relationship between a set of known 3D object coordinates (from the real-world) and their corresponding set of 2D projections on the imaging sensor (pixel coordinates).

For *camera calibration*, a simple approach to establish the 2D to 3D correspondences is by capturing a printed calibration board composed of checkers, similar to the one shown in Figure 2.4. Since the calibration board squares have known physical dimensions, their corners

Chapter 2. Simultaneous geometric and radiometric calibration of a projector-camera pair

provide us with the 3D object coordinates. After we detect the same corners in the captured images, i.e., we identify their 2D image coordinates, we can calibrate the camera with the objective of establishing a link between the two sets of coordinates. More formally:

$$s\mathbf{x}^I = \mathbf{K}[\mathbf{R}^I \mid \mathbf{t}^I]\mathbf{X} \quad (2.1)$$

where \mathbf{X} and \mathbf{x}^I are the 4D and 3D homogeneous coordinate vectors of the corners on the printed board, and their images in the captured image I , respectively. Rotation matrix \mathbf{R}^I and translation vector \mathbf{t}^I relate the optical center of the camera to the board origin (usually an extreme corner on the board). The \mathbf{K} matrix stores the intrinsic parameters of the camera, and s is an arbitrary scaling factor.

Since the board is flat, all corners of the printed pattern have equal values for the height. If we take these values to be 0, then the relationship between the points on the calibration board and the image sensor is given by a homography:

$$\mathbf{H}^I = \mathbf{K}[\mathbf{r}_1^I \mid \mathbf{r}_2^I \mid \mathbf{t}^I] \quad (2.2)$$

where \mathbf{r}_1^I and \mathbf{r}_2^I are the 1-st and the 2-nd columns of \mathbf{R}^I .

The model given in Eq. (2.1) is linear and does not take into account the possible radial and tangential distortions imposed by the camera lens. However, they also need to be estimated for an accurate camera calibration. For more details about camera calibration with flat boards, see [106].

The *geometric calibration* of the projector is slightly more complicated, as it cannot capture 3D coordinates from the real-world. In order to calibrate a projector, we use a calibrated camera, which captures the corners of a projected calibration pattern. Since the camera is precalibrated, we can use it to back-project rays through the captured projected corners into the 3D world. For accurate projector calibration it is important to account for the non-linear distortions of the camera lens when back-projecting.

To solve for the scale ambiguity of the back-projected rays, in addition to the projected calibration pattern, we also need a printed calibration pattern of known dimensions that rests on the same plane where we observe the projected pattern. Since the square dimensions of the printed pattern are known, we can calculate the rotation and translation between the board's origin and the camera optical center, i.e., we can identify the board plane in space. The calibration board is defined by its normal, given by \mathbf{r}_3^I , which can be computed from the known \mathbf{K} and \mathbf{H}^I [106], and a point on that plane, which is the vector \mathbf{t}^I . By intersecting the projected rays with the plane, we obtain the 3D coordinates of the projected corners. Acquiring the 2D pixel coordinates for the projector is straightforward, since we already have a noiseless version of the calibration pattern that we project.

Radiometric calibration establishes an input to output intensity response of a display or a capture device. Off-the-shelf display devices (monitors and projectors) rarely have the option of projecting linear images, thus we observe a non-linear output response. The most commonly observed response is a power function, also called a gamma function. This is mostly due to conformity with the sRGB standard [4], whose overall gamma is approximately $\gamma = 2.2$. However, we do not radiometrically calibrate the camera [18, 58, 67], as most consumer cameras can capture linear images ($\gamma = 1$). Note that deviations from linearity can appear towards the extreme ends of saturation, thus proper camera settings should be used when capturing the images. In this work, we focus solely on estimating the gamma of the projector.

The algorithm that we describe below is fully automatic except when the operator has to physically change the orientation of the board that carries the printed calibration pattern.

2.5.1 Image capture

For programatically capturing images we use the gPhoto2 toolbox [33], as it provides complete control over the camera settings. It is also very general and supports almost all models from major camera manufacturers.

2.5.2 Corner detection

Given the colorimetric properties of our printed and projected patterns (described in Section 2.4), we use the blue channel of the captured images to extract the printed corners. Figure 2.5b shows the blue channel for the image in Figure 2.5a. As expected, the printed pattern is completely isolated from the projected pattern. These favorable conditions allow us to use an automatic checkerboard detector [26]. The detected corners are marked with red dots on Figure 2.5b.

Figure 2.5c shows the red channel of the image in Figure 2.5a. As expected, the corner contrast suffers from the specific design of the projected pattern, especially in the darker regions. This is mainly due to the gamma of the projector which compresses the dark tones, and extends the bright. Also, a faint ghost of the printed pattern can be seen, since the yellow ink is slightly absorbing in the red wavelengths.

Analogous to the printed corners, for detecting the projected corners in step 1 of the algorithm we also utilize automatic checkerboard detector. In step 2, since the prewarp forces the projected pattern corners to land between the corners of the printed pattern, we implemented a search around the expected locations of the projected corners. We use the Harris corner detector [42] to find the projected corners with subpixel accuracy.

Chapter 2. Simultaneous geometric and radiometric calibration of a projector-camera pair

2.5.3 Prewarp

Having extracted the corners of the printed and the projected patterns, we can compute homographies that relate both the camera to the physical printed/projection board, \mathbf{H}_{cb} , and the camera to the projector image, \mathbf{H}_{cp} . Building a composition of these homographies provides us with a homography that relates the printed pattern to its current location on the projector image:

$$\mathbf{H}_{bp} = \mathbf{H}_{cb}^{-1} \cdot \mathbf{H}_{cp} = \mathbf{H}_{bc} \cdot \mathbf{H}_{cp} \quad (2.3)$$

We use this homography to compute a prewarp of the projected image such that when projected, it lands right on top of the printed pattern with a shift of half a printed square. We perform the shifting by subtracting the distance of half a printed square from the 3D positions of the board corners when estimating \mathbf{H}_{cb} .

2.5.4 Geometric calibration

To perform the actual geometric calibration we use Bouguet's camera calibration toolbox [13], which implements Zhang's [106] algorithm. In addition to the intrinsic parameters, the algorithm also provides us with the extrinsic parameters that link the camera position and orientation to that of the calibration board, for each calibration board orientation. The minimal number of board orientations is 2 (for a simplified camera model with rectangular pixels). For more general models with radial and tangential distortions, 10 or more orientations are recommended, also due to noise in the corner detection phase.

Once we have the camera calibrated, we project rays through the pixel coordinates of the corners of the projected images and intersect them with the corresponding plane that carries the printed pattern to compute their 3D positions in the camera coordinate frame. By using the corresponding rigid transforms we then translate those coordinates into the board coordinate frame and calibrate the projector using the same procedure we used for the camera.

2.5.5 Radiometric calibration

As discussed in section 2.4, each square of the projected pattern is encoded with a red value that varies across squares. To easily extract the intensity value of each square in the captured image, we undo the prewarp by fitting the image of the projected pattern onto the original pattern, i.e., we apply the inverse of \mathbf{H}_{cp} . Figure 2.6 shows an example rectified projected pattern.

Of interest to us is the subset of corners that emerge by the intersection of the projected corners, specifically on top of the white squares of the printed pattern. These corners are characterized by having grid indexes of different parity. Thus, we denote this set of corners as κ , and denote a particular corner as κ_{ij} .




Figure 2.6 – The projected pattern in its original size, after having been warped into position atop the printed pattern. The green markers, κ_{ij} , indicate the corners of the projected pattern that overlap with white squares on the printed one. The inset shows the values associated with an example corner $(i, j) = (6, 11)$.

Each κ_{ij} has four adjacent squares. We denote the top-left square as $\kappa_{ij}^{(1)}$ and continue clockwise, i.e., the bottom-left is $\kappa_{ij}^{(4)}$.

Each square has two values associated with it, the projected red-channel intensity and the captured one. Let us denote the projected square intensity as $P_{ij}^{(k)}$.

Due to light variation and other sources of noise, the pixel intensity of a square in the captured image is not uniform, and we therefore consider its mean. We denote the mean intensity value of square $\kappa_{ij}^{(k)}$ in the captured image as $C_{ij}^{(k)}$.

Given P and C we estimate the response function of the projector by fitting a power function onto the data:

$$\operatorname{argmin}_{a,b,\gamma} \sum_{k=1}^4 \left\| C^{(k)} - a \cdot \left[P^{(k)} \right]^\gamma - b \right\|^2 \quad (2.4)$$

We perform this non-linear least squares optimization for each image I in the calibration set, and thus obtain a γ -value per image, which we denote γ^I .

Finally, we compute the expected gamma value for all images:

$$\gamma^\star = \mathbb{E}[\gamma^I] \quad (2.5)$$

Note that our method focuses on extracting the data, which allows for any other non-linearity to be used for explaining the response of the projector. For that purpose, Eq. 2.4 needs to be updated to reflect the new model.

2.6 Experiments

We implemented our method in MATLAB. During the experiments, we captured linear RAW images, which we demosaiced in order to extract the different color channels. In the current implementation of the algorithm, the calibration is performed at the end after the desired set of orientations has been acquired. However, this can easily be modified in order to have a running update of the calibration parameters and the reprojection error, since the corner detection runs while capturing the images.

On the hardware side, we used an Acer H6502BD projector (color, DLP, 1920 x 1080 px) and a Canon T1i digital camera (color, CMOS, 4752 x 3168 px) with a Canon EF 24-105 mm f/4L IS USM lens. The physical calibration board was arranged in a configuration of 15 x 10 squares, where each square is 9 mm in size. Throughout the testing the calibration board was kept at a distance of approximately 45 cm from the camera and the projector. An overview of the calibration setup can be seen in Figure 2.1.

There is no inherent limitation on the resolution of the equipment imposed by the algorithm. Lower resolution cameras and projectors can also be used. A good practice is for the projected pattern to not be anti-aliased in its original projection (step 1), as the hard edges between the squares will increase the contrast of the projected corners.

The size of the printed calibration board should be selected such that it will always stay within the intersection of the camera and the projector fields of view. This is because the printed corners are not self identifying, and we need to be able to identify the projected corners for the radiometric calibration. When we observe the complete printed calibration board, we make sure that the projected calibration board will also be completely visible when the two patterns are aligned (in step 2).

We capture 2 images per calibration board orientation, which is on par with most other methods (see Table 2.1). However, we perform geometric *and* radiometric calibration at the same time, whereas the other methods only perform geometric calibration. For these methods, adding radiometric calibration means that another calibration procedure needs to be performed. Our method is much more time efficient compared to capturing the two calibration image sets separately, which requires a change of the calibration board and possibly camera exposure settings.

Geometric calibration To evaluate the calibration accuracy, we used the mean reprojection error. It is the average Euclidean distance between the detected and the reprojected corner positions in all images of the calibration set. It is measured in pixels. We can use Equation 2.1 to express the projection of world point \mathbf{X} on an image I , which we denote with $\hat{\mathbf{x}}$. If we denote the detected corner positions with \mathbf{x} , the Euclidean distance with d , the number of corners on the calibration board with N , and the number of images in the calibration set with M , the

mean reprojection error could be expressed as:

$$\frac{1}{NM} \sum_i d(\mathbf{x}_i, \hat{\mathbf{x}}_i) \quad (2.6)$$

For the geometric calibration of the camera, we achieve a mean reprojection error of 0.25 px and a very low standard deviation of 0.03 px over multiple calibrations. The calibration of the projector was slightly less accurate, with a mean reprojection error of 0.46 px and a standard deviation of 0.08 px. This is an expected result, since the accuracy of the projector calibration depends on the accuracy of the camera calibration. In addition, due to the brightness of the projector, we set the camera aperture to f/22, whereas the projector had an aperture of f/2.56. The large aperture of the projector translates into shallow depth of field, which deteriorates the corner detection accuracy.

Table 2.1 shows our results in comparison with the state-of-the-art calibration results reported in the respective publications. Only the best calibration results are shown, even if several were reported. It can be seen that our results are on par with the current methods for geometric calibration.

Reporting the mean reprojection errors in pixels does not allow us to fairly compare the various methods. This is because the metric depends on the size of the pixels, thus devices with larger pixels will show lower errors in pixels for the same physical distances. We argue that it is better to express the mean reprojection errors in μm , by multiplying the results with the pixel size. In this case, our results would be $1.17 \mu\text{m}$ for the camera, and $2.48 \mu\text{m}$ for the projector.

Radiometric calibration As a ground truth for the radiometric calibration, we projected consecutively 11 uniform grayscale images with input intensities ranging from 0 % to 100 %, in 10 % steps, over a uniformly white screen. We then captured the projections with our camera, and fitted a gamma function to the captured data. This approach is similar to [72]. We denote this method as the *ground truth method*. We estimated the gamma on different locations across the projection screen. The computed mean gamma is 2.57, with a standard deviation

Table 2.1 – Mean reprojection errors (RMSE in pixels) for camera and projector calibration and the number of images per orientation as reported by the authors.

Authors	Camera	Projector	imgs./or.
M. Ashdown [7]	0.25	0.47	-
S. Audet [8]	0.33	0.20	2
C.-Y. Chen [15]	-	0.188	2
A. Griesser [35]	<0.4	<1.5	multiple
Z. Li [57]	0.09	0.149	12
D. Moreno [68]	0.329	0.145	≈ 40
Ours	0.25	0.46	2




Figure 2.7 – Gamma estimation for 14 image calibration set. The thick red line is the mean gamma curve. The thin lines represent the estimated gamma curves for each individual image. The vertically distributed blue dots show the input vs. output pixel intensity values over all images.

of just 0.03. The low standard deviation indicates that the projector gamma is fairly uniform across the screen.

We then performed radiometric calibration with our proposed method by capturing 14 images at different orientations. These are the images from step 2 of our algorithm, which are also used for geometric calibration. Figure 2.7 graphically summarizes our results. Each blue vertical bar represents the extracted output intensities across all captured images for a single input intensity. The thick red curve is the average power function, denoted by γ^* in Eq. 2.5. The 14 thin curves are the gamma functions estimated for each image individually. The mean gamma is 2.53, with a standard deviation of 0.15. This result is in accordance with the gamma we computed with the ground truth method.

The spread of the fitted gamma curves in Figure 2.7 is mainly due to changes in the scaling and bias parameters of the fitted functions (parameters a and b in Eq. 2.4). This happens because the captured intensities are scaled by the different orientations that the projector and the camera take with respect to the calibration board (changes in irradiance). Gamma has no effect on the extreme (0 % and 100 %) input intensities, but only on the intermediate values. Furthermore, the coefficient of variation of gamma across the images is 6.1 %, whereas for the scaling parameter it is 15.1 %.

Planarity test We performed a planarity test to validate the performance of the radiometric calibration in a real-world scenario. The test consists of fitting a plane into a point cloud and estimating their fit. Therefore, we reconstructed a flat board by using a phase-shifting SL code.




Figure 2.8 – Systematic errors in the reconstruction of a flat surface. (a) Non-gamma compensated projection patterns exhibit large depth variations (RMSE = 0.331 mm). (b) Gamma compensated patterns, gamma computed with the ground truth method (RMSE = 0.056 mm). (c) Gamma computed with our proposed method (RMSE = 0.055 mm).

This is a continuous coding method and is based on the projection of sinusoidal patterns (see detailed explanation in [87]).

We performed three reconstructions of the flat surface, of size 10 x 10 cm, by projecting: non-gamma compensated patterns, gamma compensated patterns with gamma computed using the ground truth method, and with gamma computed using our proposed method. The reconstructed point cloud for each of the methods can be seen in Figure 2.8. As expected, the non-gamma compensated patterns exhibit the largest depth variations, produced by erroneous correspondences between the linear input and the non-linear output patterns. The gamma compensated patterns, on the other hand, show a much lower degree of variation.

Finally, we computed the RMSE of fitting a plane into the point clouds. The error for the linear patterns amounted to: 0.331 mm; for gamma estimated with the ground truth method: 0.056 mm; and gamma estimated with our proposed method: 0.055 mm. The 6 fold decrease in reconstruction error strongly reinforces the need for radiometric calibration, i.e., gamma estimation and compensation.

2.7 Limitations

The presented method also has some limitations. First, we assume that the gamma is constant across the projector image and across different color channels, which might not be true for lower quality equipment. In addition, severe vignetting of the camera and of the projector can also decrease the accuracy of the radiometric calibration. Initial full illumination images can therefore be taken to compensate for vignetting before projecting the patterns. Excessively specular calibration boards can degrade the quality of the extracted radiometric values. The guided detection of the projected corners in step 2 will not work well for projectors with strong radial distortions. The radiometric calibration will also decrease in accuracy in situations where the projector has a larger dynamic range than the camera. This is mostly a hardware limitation, that can be solved with multi-exposure approaches or ND filters.

2.8 Summary

In this chapter we presented a novel method that allows for geometric calibration of a projector-camera pair, simultaneous with a radiometric calibration of the projector. The method is automatic, simple, and efficient compared to other calibration methods, since it does not have an inherent overhead of projections and captures. It is faster compared to dedicated geometric and radiometric calibration techniques as it requires less manual manipulation together with a decreased complexity of a single image capturing and processing cycle. System recalibration can thus be performed more often in order to guarantee consistent results over time. We have also shown that the method offers high quality calibration results and performs on par with the current projector-camera calibration methods.

In addition, we demonstrated that radiometric calibration is essential for achieving high depth estimation accuracy with SL systems. Due to their numerous advantages, the continuous coding techniques represent the most popular choice for scene encoding, which makes our method very attractive to the 3D reconstruction community.

3 Surface Roughness Estimation Using Structured Light Projection

Surface roughness has a strong influence on material appearance. Its accurate characterization is therefore important for creating realistic reconstructions of objects in a physical or virtual environment. It has traditionally been measured with profilometers, which can be contact or optics based. The optical profilometers commonly use white light interferometry or are confocal microscopes. Figure 3.1a shows the Altimet AltiSurf 50 profilometer, that we used in our experiments. These instruments, although very accurate, are expensive and have a fairly limited scanning area. In this chapter we thus design and evaluate a surface scanning system based on structured light (SL) projection to recover the surface structure of near-planar materials. More specifically, we are interested in measuring the surface roughness of coated printing textiles, where the weave threads form macroscopic structures, with sizes in the order of 0.1 mm to 1 mm.

The difficulty of obtaining a product appearance which is close to the expected one is a limiting factor for digital printers. Modern 2.5 D printers can correctly (re)produce intricate prints in terms of both color and surface topology (e.g., master paintings), however, differences in the ink deposition process and curing, and in the composition of the inks is generally revealed by wrong light reflections, as shown in [24, 23, 77]. Recent color prediction models for printing account for the roughness of the inks [78], but do not take into account the roughness of the substrate, which is about two orders of magnitude greater for textile substrates. Accurate surface reconstruction could allow for devising color prediction models for macroscopic-level features that can potentially incorporate local interreflections. Work in this direction has already been performed, for the more specialized, but very important case of diffuse [85] and specular V-cavities [84, 86].

This chapter is structured as follows. In Section 3.1, we discuss why we chose to use a SL setup for these measurements, and what are some of the difficulties that we need to be aware of when choosing the scene encoding strategy. Then, in Section 3.2, we review the main categories of SL encoding methods, and briefly summarize their merits and shortcomings. Section 3.3 contains the experimental part, where we discuss our measurement setup, and show the surface structure and roughness that we could attain for different samples. We also compare our scans to those of an Altimet AltiSurf 50 profilometer.


3.1 Accurate depth estimation with structured light

In an effort to overcome the high cost and the limited scanning area of the profilometers, we constructed a setup for depth estimation based on SL projection. This method provides non-contact means of depth estimation by using a calibrated camera-projector pair. The simple SL setup that we constructed is composed of off-the-shelf components, and it can be seen in Figure 3.1b. The goal is to compute an accurate and high resolution depth map of the surface structure, which can be used to compute surface roughness. The high resolution is required to capture the intricate surface features of the fabrics, and the accuracy of the reconstruction could allow for devising accurate color prediction models that incorporate local interreflections.

Care has to be taken when using structured light scanning techniques, because the measured objects can be translucent and experience significant amounts of indirect illumination [38, 39, 16]. Furthermore, the scenes can be dynamic due to object movement and vibrations [93, 12]. Given these optical and mechanical noises, many coded structured light techniques fail to decode the patterns correctly, which hinders high-precision reconstruction.

Indirect or global illumination is the set of effects that include interreflections, subsurface scattering and defocus. These effects alter the directly projected patterns. Interreflections are considered as long-range effects, since they generally appear as low-frequency waves across the captured scene. Subsurface scattering, and defocus, on the other hand, are considered as short-range effects, as their effect is local and manifests in a local low-pass filtering of the projected pattern.

Interreflections are mainly caused by concavities in the measured objects' shape, where the directly projected light gets reflected from the object surface onto other portions of the object. It is generally assumed to be a reflection with low spatial frequency, however, smooth metallic



(a) Altimet AltiSurf 50



(b) A simple SL setup

Figure 3.1 – Different instruments for recovering surface structure. (a) A profilometer based on the confocal principle. (b) Our simple SL setup used in the experiments.




Figure 3.2 – An image, taken with the Altimet AltiSurf 50, of a 5 mm × 5 mm section of Canon IJM 618 — a coated polyester textile used as a print media.

or plastic surfaces can produce high-frequency mirror reflections. Subsurface scattering happens with translucent objects, where the light entering the object at one point experiences multiple scattering events and internal reflections and exits the object over a small area around the entry point. Subsurface scattering together with camera/projector defocus act as low-pass filters upon the projected patterns.

In addition to degrading the projected patterns, subsurface scattering can introduce measurement bias. During the scanning of Michelangelo's David [56], the authors measured a depth bias of approximately $40\text{ }\mu\text{m}$ for light at normal incidence for the Carrara Statuario marble. This happened because the incident light from the laser beam got scattered and formed a volume below the surface of the marble. This created a false light centroid that was detected by the camera instead of the direct reflection from the object's surface.

Scene dynamics are another major source of artifacts that mainly affect the techniques that project multiple patterns when encoding the measured scene. For example, when measuring a part of the human body (e.g., face, eye curvature, etc.) it is highly unlikely that the subject will stay still during all the projections. Furthermore, when scanning a scene with high precision even the slightest vibrations can introduce measurement noise [12].

We are interested in estimating the surface profile of printing fabrics. They have near planar surfaces and are white, since they are used as print media. An example 5 mm x 5 mm patch of a Canon IJM 618, white, polyester textile coated with a special polymer coating, is shown in Figure 3.2.

3.2 Structured light patterns for surface profilometry

Salvi et al. in [87] present a comprehensive overview and classification of a large variety of coded structured light (CSL) techniques. The aim of CSL is to encode the position of each pixel (or a region) in the projected pattern in a unique way, such that a correct correspondence can be established between the projected and the captured patterns.

This field of research has been quite active in the past three decades. A myriad of CSL techniques is now available, all offering different trade-offs that are suited for specific purposes. [87] cites more than forty CSL methods and classifies them hierarchically. Covering all different classes of algorithms in detail will require a long discussion. Therefore, in this section we will focus on the most general algorithms, and we will skip their variations developed for specific purposes. The coarsest level of classification divides the CSL patterns into discrete and continuous coding methods.

3.2.1 Discrete Coding Methods

Discrete coding methods are characterized by the discrete nature of the projected pattern(s), i.e., the projected signal varies in discrete and distinct steps across its dimensions. In general, they use an absolute coding strategy, as opposed to a relative coding strategy, meaning that each projected stripe can be uniquely identified.

The discrete coding techniques are further divided into spatial and time multiplexing. The spatial multiplexing patterns use the surrounding of the features to encode their locations in the pattern, while the time multiplexing patterns localize the pixels by projecting multiple successive patterns, where each pattern provides a piece of the pixels' unique codewords. Therefore, the spatial multiplexing techniques require only a single projected pattern, and the time multiplexing techniques require a series of projected patterns. The trade-off between projecting single or multiple patterns is that although a single pattern can be used for imaging moving targets, to achieve the same precision they need larger alphabet sizes (e.g., number of distinct colors, or patterns), which decreases their noise resistance. Therefore the temporal multiplexing techniques tend to achieve higher spatial accuracy and density.

De Bruijn sequence is the most commonly used pattern among the spatial multiplexing methods. A k -ary De Bruijn sequence of order n is a cyclic sequence, created from an alphabet of size n , that contains each subsequence of length k only once. The length of the sequence is $n^k - 1$. This sequence is interesting for structured light coding since by using a 1D sliding window of size k along the epipolar lines, we can uniquely identify the position of each projected stripe in the captured image. An example De Bruijn sequence composed of $n = 4$ colors and window size of $k = 3$ stripes can be seen in Figure 3.3.

The measurement quality of these methods is directly linked to the number of projected stripes or slits (the stripes are adjacent to each other, while the slits are stripes separated by




Figure 3.3 – De Bruijn pattern with $n = 4$ and $k = 3$. Image taken from [87].

black gaps). The number of stripes that can be projected is equal to the length of the De Bruijn sequence, which in turn depends on the codeword basis n and the length of the subsequences k . If k increases, the algorithms have difficulties establishing the correspondences in areas where the pattern is obscured by shadow. In addition, the algorithms are more susceptible to noise, since the probability that at least one of the stripes from the sliding window will be affected by noise increases. If, on the other hand, n increases, then the radiometric distance between the codewords decreases, which increases the chance of wrongly interpreting some codewords. Common approaches use 3-6 fully saturated and maximally distant colors in the hue channel for the alphabet, and vary k in order to achieve the desired sequence length.

Among the first proposed strategies for structured light coding were the time multiplexing codes that were introduced in 1982 by Posdamer and Altschuler [80]. The idea behind the *binary* time multiplexing codes is simple: vertical binary stripes (usually black-white) are being projected subsequently, where in each successive projection the number of stripes doubles. The unique codewords are created after all the patterns have been projected. These projections usually follow a coarse-to-fine approach until the desired resolution is reached. The number of addressable positions (in a single dimension) is 2^m , where m is the number of projected patterns. The maximal achievable resolution is limited by the projector and/or the camera used. Since the codeword basis in each projection is small, this approach provides an increased resistance to noise, with the downside of not being able to capture moving scenes.

Gray code is the usual *binary* encoding of choice for time multiplexing coding. Unlike the binary code, in Gray code the consecutive codewords have a Hamming distance of one. This makes the encoding more resistant to noise, since it opens the possibility for error correction. An example sequence of ten Gray code projections can be seen in Figure 3.4.

The n -ary temporal codes go along the lines of the *binary* codes; they only increase the size of the alphabet for each projected pattern. Their alphabets commonly use colors, e.g. {red, green, blue}. The number of unique addressable positions in the projected patterns can then be expressed by n^m , where n is the alphabet size. This approach decreases the number of projections for achieving the same spatial resolution, however, it makes the patterns more prone to errors when the measured surface is colored.




Figure 3.4 – 10 successive projections that encode the scene with Gray code. The ordinal number $m = \{1, \dots, 10\}$ is shown on the left side of each pattern.

Once a feature point has been identified in the projected and the captured images, a triangulation procedure is carried out to recover the 3D position of the imaged surface point. Since the triangulation involves tracing rays in 3D space, a geometric calibration of the camera-projector pair is a prerequisite. A radiometric calibration is usually required only by the methods that use colored patterns, since they have to cope with changing surface albedos.

3.2.2 Continuous Coding Methods

The continuous coding methods are characterized by a continual variation of the coding intensities along the coding axes. Grayscale patterns are the standard choice for these methods, which make them effective on object surfaces of any color. The methods in this group employ periodic encoding.

The phase shifting methods are a subclass of the continuous coding methods. When projecting a periodic pattern on a surface, each point can be characterized by its phase offset in the pattern. If there are non-flat objects in the scene, they will cause a phase deviation in the pattern when compared to the same pattern projected on a flat reference surface (plane). This deviation can be used to recover the height of surface points relative to the reference plane. The projected signal can be expressed as:

$$I_n^p(x^p, y^p) = A^p + B^p \cos(2\pi f_\Phi x^p - 2\pi n/N), \quad (3.1)$$

where x is along the coding axis, the superscript p signifies the projected pattern, A^p and B^p are the signal offset and amplitude, f_Φ is the carrier frequency, (x^p, y^p) are the projection pattern coordinates, and finally $n = 0 \dots N - 1$. The pattern that the camera captures is the following:

$$I_n(x, y) = \alpha(x, y)[A^c + B^c \cos(2\pi f_\Phi x^p + \phi(x, y) - 2\pi n/N)] \quad (3.2)$$

The captured pattern experiences both an intensity change and a phase shift. The phase deviation is extracted as follows:

$$\phi(x, y) = \arctan \left[\frac{\sum_{n=1}^N I_n(x, y) \sin(2\pi n/N)}{\sum_{n=1}^N I_n(x, y) \cos(2\pi n/N)} \right] \quad (3.3)$$

The retrieved phases are independent of the surface albedo. The minimal number of captured shifted patterns is $N = 3$, since there are three unknowns in the system: A^c , B^c and $\phi(x, y)$ [11]. Higher carrier frequencies yield higher signal-to-noise ratio, but are also more sensitive to sharp object edges.

With this computation we only recover the wrapped phase of the surface points. It is called wrapped since the arctan function wraps the values in the range $(-\pi, \pi]$. Therefore a phase unwrapping algorithm has to be used to recover the natural phases of the pixels. A simple approach to phase unwrapping would be summing the gradients of the wrapped phases along a continuous path, starting from a reference point and going to all other surface points. This approach assumes that the absolute values of the gradient are always less than π , and if there is a π to $-\pi$ discontinuity, it adds 2π (or vice versa). The height of a surface point relative to the reference plane is then calculated as follows:

$$h(x, y) = l_0 \frac{\hat{\phi} - \hat{\phi}_r}{\hat{\phi} - \hat{\phi}_r - 2\pi f_\Phi d}, \quad (3.4)$$

where $\hat{\phi}$ and $\hat{\phi}_r$ are the unwrapped phases imaged at the same pixel location (x, y) for the measured surface and the reference plane, respectively, l_0 is the distance from the projector and camera lenses to the reference plane, and d is the baseline distance between the projector and camera lenses.

An obvious drawback of this method is the necessity to project multiple patterns. Frequency multiplexing is another large subclass of continuous coding techniques. The difference with the phase shifting methods is that they perform the phase decoding in the frequency domain. They have an advantage over the phase shifting techniques because they require a single projection to extract the height, thus they can deal with dynamic scenes.

Similar to the phase shifting methods, the height of the object is encoded in the phase deviation. Therefore, it needs to be separated from the unwanted amplitude variations in the captured image. In Fourier Transform Profilometry (FTP), a 1-dimensional Fourier transform is performed on each row of the captured image, that we denote with $G(f, y)$. An example magnitude spectrum of $G(f, y)$ can be seen on Figure 3.5.

Due to the periodic nature of the projected pattern, there is a maximal slope that can be resolved. Greater slopes can interfere with the fundamental component of the signal. This can be visualized on Figure 3.5 with $(Q_b)_{max}$ crossing over $(Q_1)_{min}$.




Figure 3.5 – Frequency magnitude spectrum of the captured image for FTP. The height information is stored in the Q_1 spectrum, and it needs to be isolated through band-pass filtering.

To extract the phase deviation from the captured images between the reference pattern and the captured pattern, a band-pass filtering is performed and only the spectrum denoted by Q_1 on Figure 3.5 is selected, which corresponds to the projected signal. An inverse Fourier transform then renders an image of the clean signal. Since the extracted phase is wrapped in the $(-\pi, \pi]$ interval, a phase unwrapping is necessary to recover the natural phase.

3.3 Experiments

To test the accuracy of our SL setup, we reconstructed the surface structure of printing textiles. We compared our reconstruction qualitatively and quantitatively, by using the surface roughness measure, with those of a profilometers.

An example $5 \text{ mm} \times 5 \text{ mm}$ patch of a Canon IJM 618, a white, polyester textile coated with a polymer coating, is shown in Figure 3.6a. As it can be seen, the measured fabrics are near planar surfaces, and are white since they are used a print media. Table 3.1 summarizes the main characteristics of the measured textiles.

Table 3.1 – The characteristics of the textile samples that we used for testing.

Supplier	Product name	Weight [g/m^2]	Thickness [mm]	Fabric type
Canon	IJM 618	200	0.27	Woven
Canon	IJM 617	310	0.39	Woven
Senfa	Decoprint Night	320	0.33	Woven
Senfa	Decoprint Grain	375	0.45	Woven
JM	Mediatex Twintex FR	380	0.61	Knitted




Figure 3.6 – (a) An example printing textile that we scanned. The measured fabrics are white near planar surfaces, since they are used as a print media. (b) The structured-light setup used for the experiments

3.3.1 Structured light encoding strategy

We should consider the indirect illumination effects when scanning these textiles to choose the best encoding strategy. The local interreflections should not cause major issues when decoding the projected patterns, since the surface features are relatively small and cannot reflect considerable amounts of light towards other parts of the fabric. Subsurface scattering is, however, present with these materials. This will result in smoothing of the projected patterns, and can cause problems with high-frequency binary patterns. Finally, these are static scenes, and allow us to use temporal encoding techniques that project multiple patterns.

Taking these considerations into account, we decided to use a continuous phase shifting encoding. Unlike the discrete coding methods, these method allow us to use the full resolution of the camera, which is higher than that of the projector. The slight blurring of the patterns due to sub-surface scattering does not affect these patterns. Furthermore, it is often suggested to slightly defocus the projector, such that the pixel grid will not be visible in the camera image.

3.3.2 Measurement setup

We used a very simple structured light setup composed of a single camera and a single projector, which is shown on Figure 3.6b. The projector that we used was an Acer H6502BD (color, DLP, $1920 \text{ px} \times 1080 \text{ px}$) and the camera was a Canon 5D Mark II (color, CMOS, $5616 \text{ px} \times 3744 \text{ px}$) with a Canon EF 50 mm f/2.5 Macro lens. The samples were affixed to the flat board shown in the background.

To calibrate the camera and the projector geometrically we used the simultaneous geometric-radiometric calibration method described in Chapter 2. The camera is capable of capturing linear images, however, the projector applies non-linear processing to the input signal. Since




Figure 3.7 – The measurement principle for our SL setup. The measured scene is illuminated with structured light by the projector. Since we use a continuous coding technique, phase shifting, and a non-linear projector, the values of the projected patterns need to be distorted in the opposite direction, such that when projected they become linear. The projected patterns are captured by the camera, and a correspondence is established between the projector and camera pixels. A triangulation procedure then computes the 3D position of each camera pixel, for which a geometric calibration of the projector-camera pair is required. The output of this process is a point cloud, containing a point for each camera pixel.

we are using a continuous coding technique, we performed a radiometric calibration of the projector. Instead of capturing a separate set of images for the radiometric calibration, we received it for free by using the above calibration method.

3.3.3 Measurement principle

The measurement principle of our setup is shown in Figure 3.7. The measured scene is illuminated with structured light by the projector. Since we are using a continuous coding technique (phase shifting), and a non-linear projector, the values of the projected patterns need to be distorted by an equal amount, but in the opposite direction. When projected, they become linear, which is what the decoding algorithm expects in the captured images. The projected patterns are captured by the camera, and a correspondence is established between the projector and the camera pixels. With the continuous coding techniques, we are limited by the resolution of the camera, because we can establish sub-pixel correspondences with the projector image. A triangulation procedure then computes the 3D positions of each camera pixel, for which the geometric calibration of the projector-camera pair is required. The output of this process is a point cloud, with as many points as there are camera pixels.

Given the periodic nature of the projected pattern, we chose to unwrap its phase with projections of sinusoidal patterns that have progressively smaller wavelengths. We decided to use this approach, since the algorithms for phase unwrapping have difficulties in the presence of noise, local shadows, irregular surface brightness, etc, and we are interested in an accurate depth estimation. Thus, we projected 24 patterns: 8 wavelengths, and 3 phase shifts of 120° for each wavelength. We started with a pattern with 2048 px wavelength, which covers the projector image with a single period, and we progressively halved it until a pattern with a 16 px wavelength, which was the smallest wavelength that the projector could project reliably.

3.3.4 Ground truth measurements

Ground truth measurements were performed with an Altimet AltiSurf 50 profilometer, which can be seen in Figure 3.1a. It is based on the confocal principle. It illuminates a single point of the object with, initially a white light, that gets spread into different wavelengths by a lens through axial chromatic aberration. These different wavelengths cover a range of depths. The reflected light from the object then passes through an aperture, and is registered by a CCD spectrometer. The wavelength where the signal is the strongest is directly related to the depth of the measured sample.

The surface structure of the Canon IJM 618 sample from Figure 3.6a, measured with the Altimet AltiSurf 50 profilometer, can be seen in the top figure of Figure 3.13. The x-, y- and z-axis steps were $10\text{ }\mu\text{m}$, $10\text{ }\mu\text{m}$ and $0.156\text{ }\mu\text{m}$, respectively. The measured area was $5\text{ mm} \times 5\text{ mm}$, yielding 501×501 data points. In the figure we can see that the resolution in every axis is very fine, such that we can see the fibers that make up the threads. We can also see that there are very few erroneous pixels, however, some anomalies in the valleys surrounded by four threads are visible. Those can either be due to the measuring process or they can simply exist due to the coating or weaving of the material. That is already a very fine scale, and we are only interested in the roughness measure of the surface. The measurements with the Altimet AltiSurf 50 were provided to us by Océ with the help of Marine Page.

3.3.5 Roughness metric

Roughness can be defined in several ways, and here we will use the root mean square height of the surface points with respect to the mean plane of the surface they comprise [48]:

$$S_q = \sqrt{\frac{1}{A} \iint z^2(x, y) dx dy}, \quad (3.5)$$

where A is the area of the measured surface, z is the height from the mean plane of a data point with x and y its horizontal and vertical coordinates, respectively. In the case of a uniformly discretized surface, as in our case, we can write:

$$S_q = \sqrt{\frac{1}{N} \sum_i \sum_j z^2(x_i, y_j)}$$

To test the geometric calibration of the projector-camera pair and the radiometric calibration of the projector, we reconstructed a portion of the flat board that holds the samples. The reconstruction of the board after a careful calibration with 15 patterns is shown in Figure 3.10.

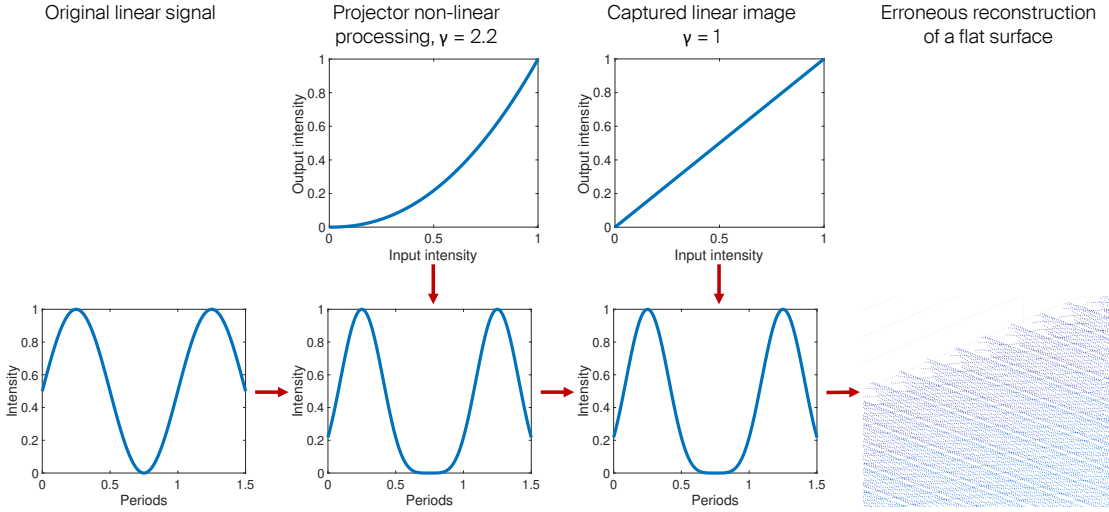


Figure 3.8 – Source of periodic surface reconstruction errors due to non-linear processing of input signals. The radiometric calibration is important when the camera or the projector introduce non-linear processing to the input images, and the structured light encoding method features continuous intensity variations. For example, the phase shifting scene encoding depends on projecting sinusoidal patterns. The projector applies non-linear processing to the input signal. Although the camera captures linear images, the phase shifts in the processed signal lead to erroneous correspondences between the camera image and the original pattern. This leads to periodic errors in the surface reconstruction.

The top figure shows a 5 mm × 5 mm patch, corresponding to the size of the AltiSurf 50 scan. We can see that the reconstruction is fairly accurate, with a roughness of only 3.24 μm. The visible waves that propagate along the x-axis in the height profile are caused by errors in the radiometric calibration.

3.3.6 Reconstruction errors due to non-linear processing

In Figure 3.8, we can see the periodic errors that appear when the projector or the camera apply non-linear processing to input patterns that feature continuous intensity variations, such as in the phase shifting encoding. The original signal that we intend to project is a perfect sine wave. The projector then applies a γ of 2.2, which distorts the sine wave. The slopes become steeper and the same signal intensity can now be found either earlier (after a peak), or later (after a valley), compared to the original sine wave. We can also view this as phase shifting the original signal positively or negatively. If we then project this processed signal onto a flat surface, and establish correspondences between the intensity values, or the phase information, between the original and processed signals, we would recover the phase shifts. These phase shifts would normally appear because of depth variations in the scene, which is why we see errors in the reconstruction of the flat surface. It is therefore necessary to radiometrically calibrate the device that applies non-linear processing, in this case the projector. Once we know the non-linear response of the projector, we can apply the opposite function to the original patterns, such that when projected, we get the intended sine signal.




Figure 3.9 – Global fabric deformation for the 5 mm \times 5 mm Senfa Decoprint Grain scan. The colormap on the right-hand side color codes the height deviation from the best fitting plane, and it is expressed in μm . (a) The fabric deformation for a coarse 3×3 grid. Each cell represents the average deviation of the data points inside it. (b) The compensation that we will subtract from each data point of the point cloud. (c) The average height deviation of each cell after the compensation has been applied.

3.3.7 Error compensation

The fabrics are flexible, and might bend slightly even when attached to the flat supporting surface. Therefore, we performed a global rectification of all samples, measured with both the Altisurf 50 and with our setup. Figure 3.9 shows the deformation and compensation for an example scan. We divided the sample's surface into a coarse 3×3 grid, and calculated the average height from the best fitting plane for each cell (Figure 3.9a). We then applied an interpolation to compute a compensation for every point of the sample's point cloud (Figure 3.9b). The coarse grid should capture the global distortion of the fabric from the best fitting plane, if there is any. The expectation is that the coarse grid will capture an area where the average deviation from the best fitting plane should be zero, since we expect the samples to be near-planar, and with repetitive surface features. A denser subdivision of the patch will start to conform to the height of individual surface features, thus deviating from the expected zero mean height deviation from the best fitting plane. This compensation would not be necessary if the samples were rigid.

To eliminate the waviness due to the errors in the radiometric calibration (clearly visible in the top figure of Figure 3.10), we computed the average height of each column along the y-axis, and compensated the height profile accordingly. This is possible because our projected patterns' encoding is along the x-axis, thus the error patterns propagate exclusively along the x-axis. This error compensation assumes that the mean deviation from the best fitting plane of a column of points along the y-axis is zero, and a possible height bias might exist due to the errors of the radiometric calibration. More concretely for our measured samples, the compensation assumes that the weaving patterns are periodic, and that there are enough periods along the y-axis in the fabric that the mean deviation is zero. This claim appeared to be indeed correct, because the computed compensations, a signal formed by the mean value

across the y-axis for each x-coordinate, of all samples had very similar frequencies, and the compensation was similar to the one for the flat surface.

The bottom figure in Figure 3.10, shows the effects when both compensations are applied. This is a completely flat surface without visible macro structures on its surface. The surface roughness of $1.61\text{ }\mu\text{m}$ is therefore due to a non-planarity of the surface that was not accounted for, with the simple compensation scheme. The noise in the measurements are zero mean distributed about the measured surface.

3.3.8 Results

After we established that the calibration is correct, and that the compensations works well, we can measure the roughness of the textile samples. The surface scans for all samples from Table 3.1 are shown in Figures 3.11-3.15. The scans with the AltiSurf 50 are shown in the top subfigure of each figure, and from our setup in the bottom subfigure. The height values are computed from the best fitting plane, and are therefore both positive and negative. The resolution of our scans is much lower in all 3 dimensions. The point clouds contain around 11'750 points for a grid of approximately 107×107 , compared to the 251'001 points in those of the AltiSurf 50, which yields a step of approximately $46.7\text{ }\mu\text{m}$ in the x- and y-axes. We computed the roughness values on the native point clouds, and we then upsampled our point clouds to the same 501×501 point grid for display. The samples were at a distance of approximately 45 cm from the projector-camera pair. It was not possible to bring the samples closer, because we were limited by the minimal focal distances and the fields of view of the projector and camera lenses.

In Figure 3.11 we can see the surface of the Senfa Decoprint Grain. Both scans reveal the structure of the weave with distinct zig-zag peaks. The period of the weave is slightly approximately 1 mm. Our scan is rougher and somewhat noisier, and registers a slightly higher roughness value of $13.59\text{ }\mu\text{m}$, as compared to $13.13\text{ }\mu\text{m}$. The difference can be caused by several factors, including slight warping of the sample, non-equal surface structure due to the weave or coating, different number of weave periods due to limited scan areas, or simply errors in the measurement. Nevertheless, the roughness values are very similar, with a difference of 3.5 %.

The surface structure of Canon IJM 617 can be seen in Figure 3.12. Here we can again see a similarity in the weave structure, with certain rows along the x-axis protruding higher. Additionally, we can see plenty of non-regular structures as well, which is mostly likely the result of the surface coating. The period of the weave is approximately 0.55 mm. The roughness values are $13.44\text{ }\mu\text{m}$ for our measurement, and $14.65\text{ }\mu\text{m}$ for the AltiSurf, for a difference of -8.22 %.

The Canon IJM 618 sample shows a very regular weave, which can be seen in Figure 3.13. Here we can see the much higher resolution of the AltiSurf. The structure is correct in our

Table 3.2 – A summary of the roughness results. The roughness computed from the 5 mm × 5 mm point clouds computed with the Altimet AltiSurf 50 and with our structured light setup are compared, and the difference is shown in the *Altimet diff.* column. The last column of the table compares the surface roughness computed from the 5 mm × 5 mm and the 10 mm × 10 mm surface scans performed with our setup. The units of the surface roughness S_q and the maximal height of the surface S_z are μm .

Product name	Altimet AltiSurf 50		SL setup		SL setup	
	5 mm × 5 mm		5 mm × 5 mm		10 mm × 10 mm	
	S_z	S_q	S_q	Altimet diff.	S_q	5 mm × 5 mm diff.
Senfa Decoprint Grain	122	13.13	13.59	3.5 %	14.46	6.41 %
Canon IJM 617	144	14.65	13.44	-8.22 %	13.99	4.07 %
Canon IJM 618	81	13.3	14.31	7.67 %	15.48	8.13 %
Senfa Decoprint Night	100	15.98	12.63	-20.97 %	13.12	3.9 %
JM Mediatex Twintex FR	247	39.87	29.46	-26.1 %	32.25	9.48 %

scan, as both the higher and the lower threads can be distinguished as yellow and faint green peaks, respectively. Here we also see some anomalies in the valleys where four threads meet, likely caused by shadowing, specular reflections or interreflections. The roughness that we compute from our scan is 14.31 μm , and from the AltiSurf scan it is 13.3 μm . This difference is likely caused by the anomalies in our reconstruction, and it amounts to 7.67 %. The period of the weave is approximately 0.6 mm, similar to the one in the Canon IJM 617, but with more pronounced peaks.

The Senfa Decoprint Night has the highest frequency weave of all textiles, with a period of 0.42 mm, and it can be seen in Figure 3.14. Our setup manages to partially capture the surface structure, but it is inconsistent across the surface. The frequency of the weave is correct, but the reconstruction does not reveal the shape of the threads, with the crevasses in-between, as in the AltiSurf reconstruction. The roughness deviation of our scan from the one of the AltiSurf is -20.97 %.

Finally, Figure 3.15 shows the surface structure of the JM Mediatex Twintex FR. This textile is knitted, and features the roughest surface, with the greatest height range. The two scans now show obvious differences. Our scan manages to recover the correct global structure, and the correct structure frequency, but fails to recover the finer structures that comprise each loop. The period of the textile is approximately 0.625 mm in the x-axis, and 0.9 mm in the y-axis. The deviation between the roughness values amounts to -26.1 %.

Figures 3.16, 3.17, and 3.18 show larger, 10 mm × 10 mm, patches of our scans of the corresponding textiles. We can see that, compared to the smaller 5 mm × 5 mm scans, the roughness increases by 6.41 %, 4.07 %, 8.13 %, 3.9 %, 9.48 %, respectively. This can be caused by the differences in the weave pattern across the surface, but, a more likely explanation is that it is caused by the non-planarity of the samples in this larger area, since it is hard to affix them completely flush with the board. A summary of all results is shown in Table 3.2.

3.4 Summary

The accurate characterization of surface roughness is important for creating realistic reconstructions of real objects and for simulating real materials in virtual environments. While surface roughness is generally measured by optical profilometers, they are expensive instruments and have a fairly limited scanning area. In this work we evaluated the feasibility of an inexpensive SL setup to be used for computing reliable surface structure and roughness. It can be used for scanning much wider surface areas, relative to those from a profilometer.

We presented the challenges that might appear when estimating surface depth with SL, due to indirect illumination and object movement. We then provided an overview of the general classes of SL encoding, and briefly discussed their positive and negative aspects.

To test the accuracy of the surface reconstruction, we selected a representative set of five printing textiles. These textiles were characterized by different weave frequencies and feature heights. The surface features of the scanned textiles had periods ranging from 0.42 mm to 1 mm, and maximum heights of the surface from 81 to 247 μm .

For comparison, we used the Altimet AltiSurf 50 profilometer, which is based on the confocal principle. After a careful calibration of our structured light setup, we managed to achieve comparable roughness values to those of the profilometer, with 3.5 % to 8.22 % difference for the materials that we could scan reliably. Such errors could arise from global warping in the shape of the samples, as these are flexible materials, and from uneven weaving and coating. The materials that we could scan reliably were the Senfa Decoprint Grain, Canon IJM 617, and Canon IJM 618, which had weave periods of 1 mm, 0.55 mm, and 0.6 mm, and a maximum surface heights of 122 μm , 144 μm and 81 μm , respectively. Some unexpected very narrow valleys were also visible in the reconstruction of the Canon IJM 618, which we suspect appeared due to interreflections, specular reflections or shadowing. Some specular reflections can be seen in the image of the surface in Figure 3.7.

We also showed the limitations of our structured-light setup, which appeared when scanning the Senfa Decoprint Night and the JM Mediatex Twintex FR textiles. The Senfa Decoprint Night is a textile that has a high frequency weave (period of 0.42 mm), with deep crevasses between the threads (maximal surface height of 100 μm). The JM Mediatex Twintex FR textile on the other hand has a lower frequency knitted pattern (period of 0.625 mm in the x-axis, and 0.9 mm in the y-axis), but very prominent surface features (maximal surface height of 247 μm). For these materials we generally need higher resolution scans in order to capture the finer surface details. This can be achieved with a higher focal length lens for the camera, and optionally for the projector to gain better signal-to-noise levels.

We also showed scans of the printing textiles with larger 10 mm \times 10 mm size. The surface features of the scanned textiles were recovered with the same accuracy as in the case of the smaller 5 mm \times 5 mm scans. The computed roughness was also inline with that from the smaller scans, but slightly higher, due to the unavoidable warping of the textile samples. A table with suctioning support might be able to remedy this problem.




Figure 3.10 – Surface structure of the flat board used for affixing the samples, measured with our structured light setup. The colormap on the right side color codes the height of the surface points. The height is represented as the deviation from the best fitting plane to the scanned point cloud. All units are mm, except in the title, where they are μm . The top subfigure shows the point cloud computed by our simple structured light setup, upsampled to the resolution of the Altimet AltiSurf 50. The wavy patterns are height deviations due to errors in the radiometric calibration of the projector. The bottom subfigure shows the same point cloud, after we apply the compensation for the non-planarity of the sample and for the errors due to the radiometric calibration.

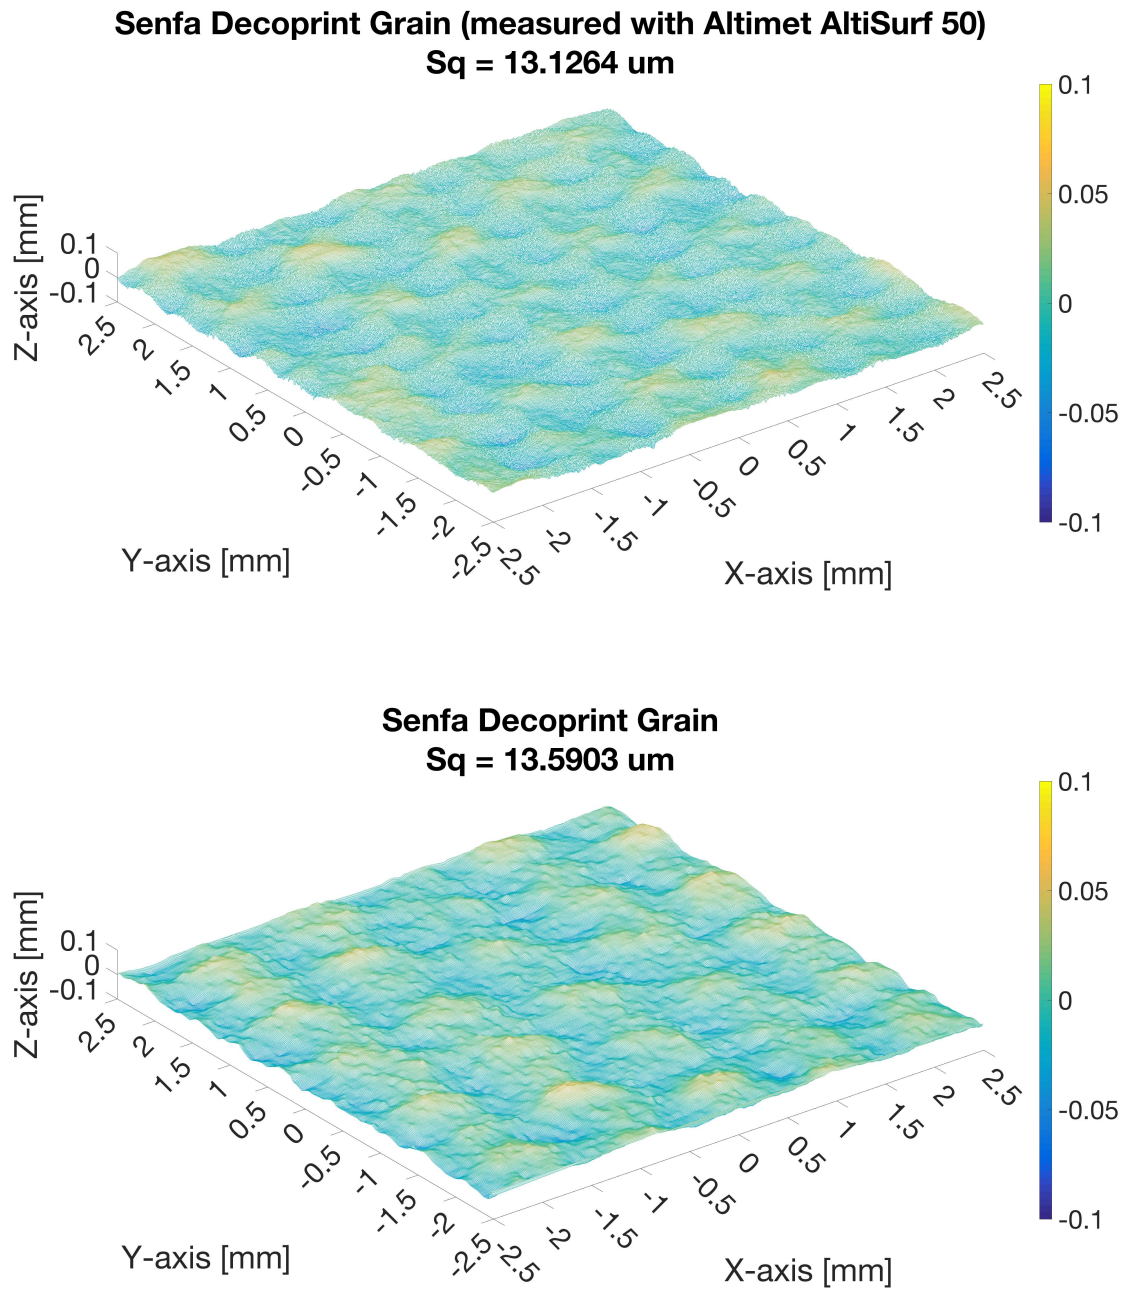


Figure 3.11 – Surface structure of Senfa Decoprint Grain, polyester textile with acrylic coating. The colormap on the right side color codes the height of the surface points. The height is represented as the deviation from the best fitting plane to the scanned point cloud. All units are mm, except in the title, where they are μm . The top subfigure shows the point cloud computed by the Altimet AltiSurf 50, after we apply our simple compensation for the non-planarity of the sample. The bottom subfigure shows the point cloud computed with our structured light setup, after we apply the compensation for non-planarity of the sample and for the errors due to the radiometric calibration. Both point clouds show a similar surface structure of the sample, and also very similar surface roughness.




Figure 3.12 – Surface structure of Canon IJM 617, polyester textile with polymer coating. The colormap on the right side color codes the height of the surface points. The height is represented as the deviation from the best fitting plane to the scanned point cloud. All units are mm, except in the title, where they are μm . The top subfigure shows the point cloud computed by the Altimet AltiSurf 50, after we apply our simple compensation for the non-planarity of the sample. The bottom subfigure shows the point cloud computed with our structured light setup, after we apply the compensation for non-planarity of the sample and for the errors due to the radiometric calibration. Both point clouds show a similar surface structure of the sample, and also similar surface roughness.




Figure 3.13 – Surface structure of Canon IJM 618, polyester textile with polymer coating. The colormap on the right side color codes the height of the surface points. The height is represented as the deviation from the best fitting plane to the scanned point cloud. All units are mm, except in the title, where they are μm . The top subfigure shows the point cloud computed by the Altimet AltiSurf 50, after we apply our simple compensation for the non-planarity of the sample. The bottom subfigure shows the point cloud computed with our structured light setup, after we apply the compensation for non-planarity of the sample and for the errors due to the radiometric calibration. Both point clouds show a similar surface structure of the sample, and also similar surface roughness. Our scan also shows anomalies in the form of valleys at irregular locations right next to the thread peaks.

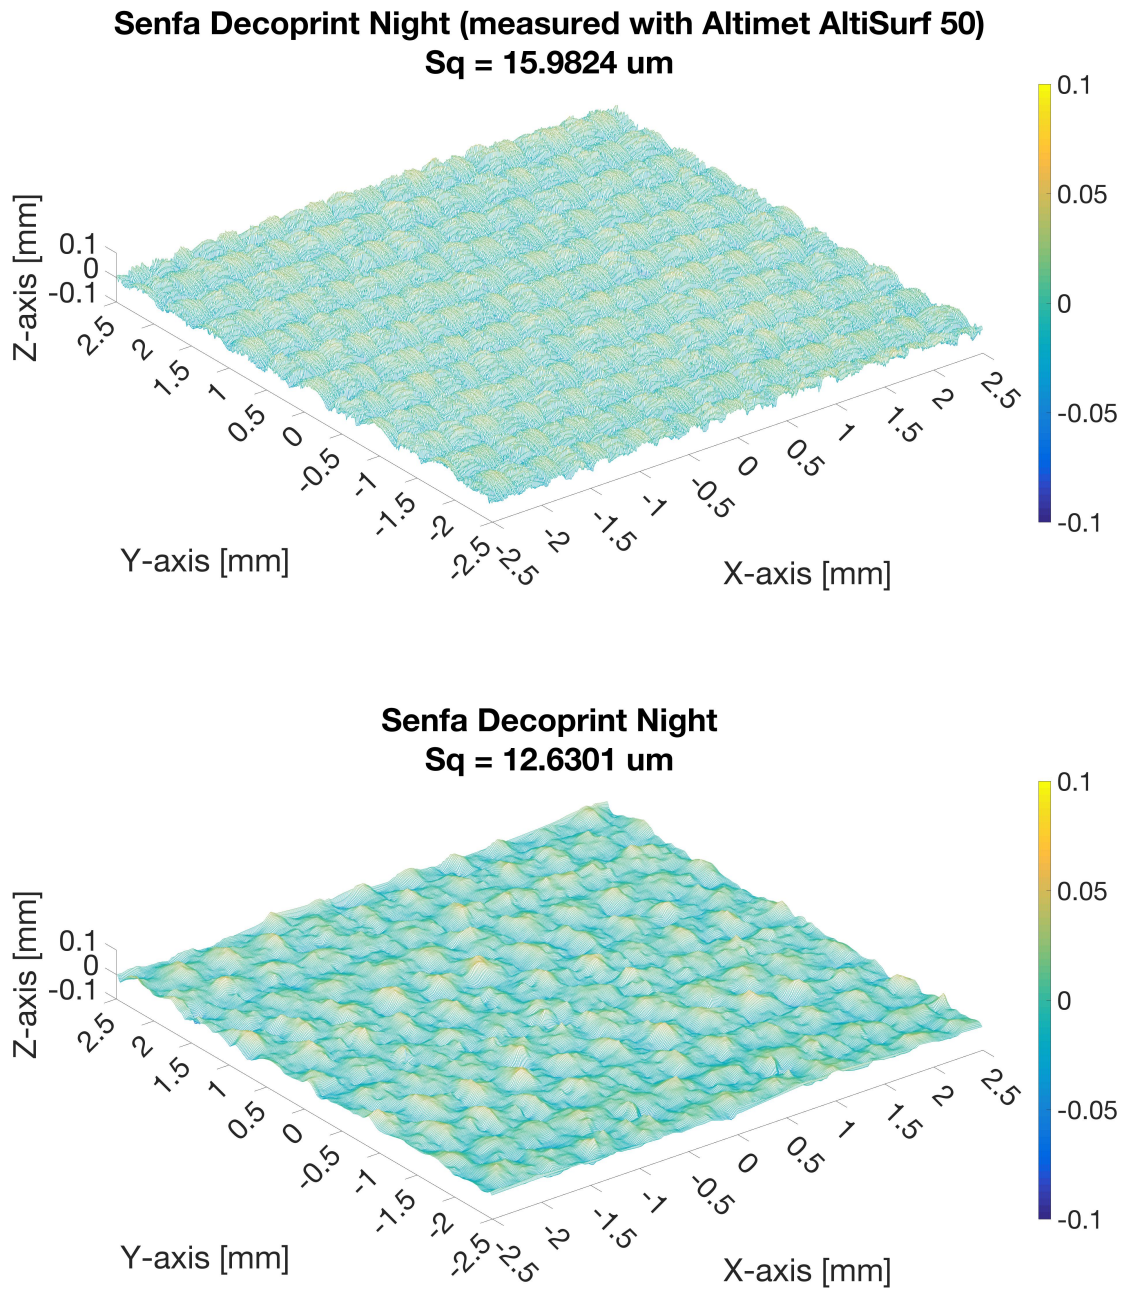


Figure 3.14 – Surface structure of Senfa Decoprint Night, polyester textile with acrylic coating. The colormap on the right side color codes the height of the surface points. The height is represented as the deviation from the best fitting plane to the scanned point cloud. All units are mm, except in the title, where they are μm . The top subfigure shows the point cloud computed by the Altimet AltiSurf 50, after we apply our simple compensation for the non-planarity of the sample. The bottom subfigure shows the point cloud computed with our structured light setup, after we apply the compensation for non-planarity of the sample and for the errors due to the radiometric calibration. Although the surface structure is similar, our scan also shows lower surface roughness values since cannot capture the high-frequency details of the weave.

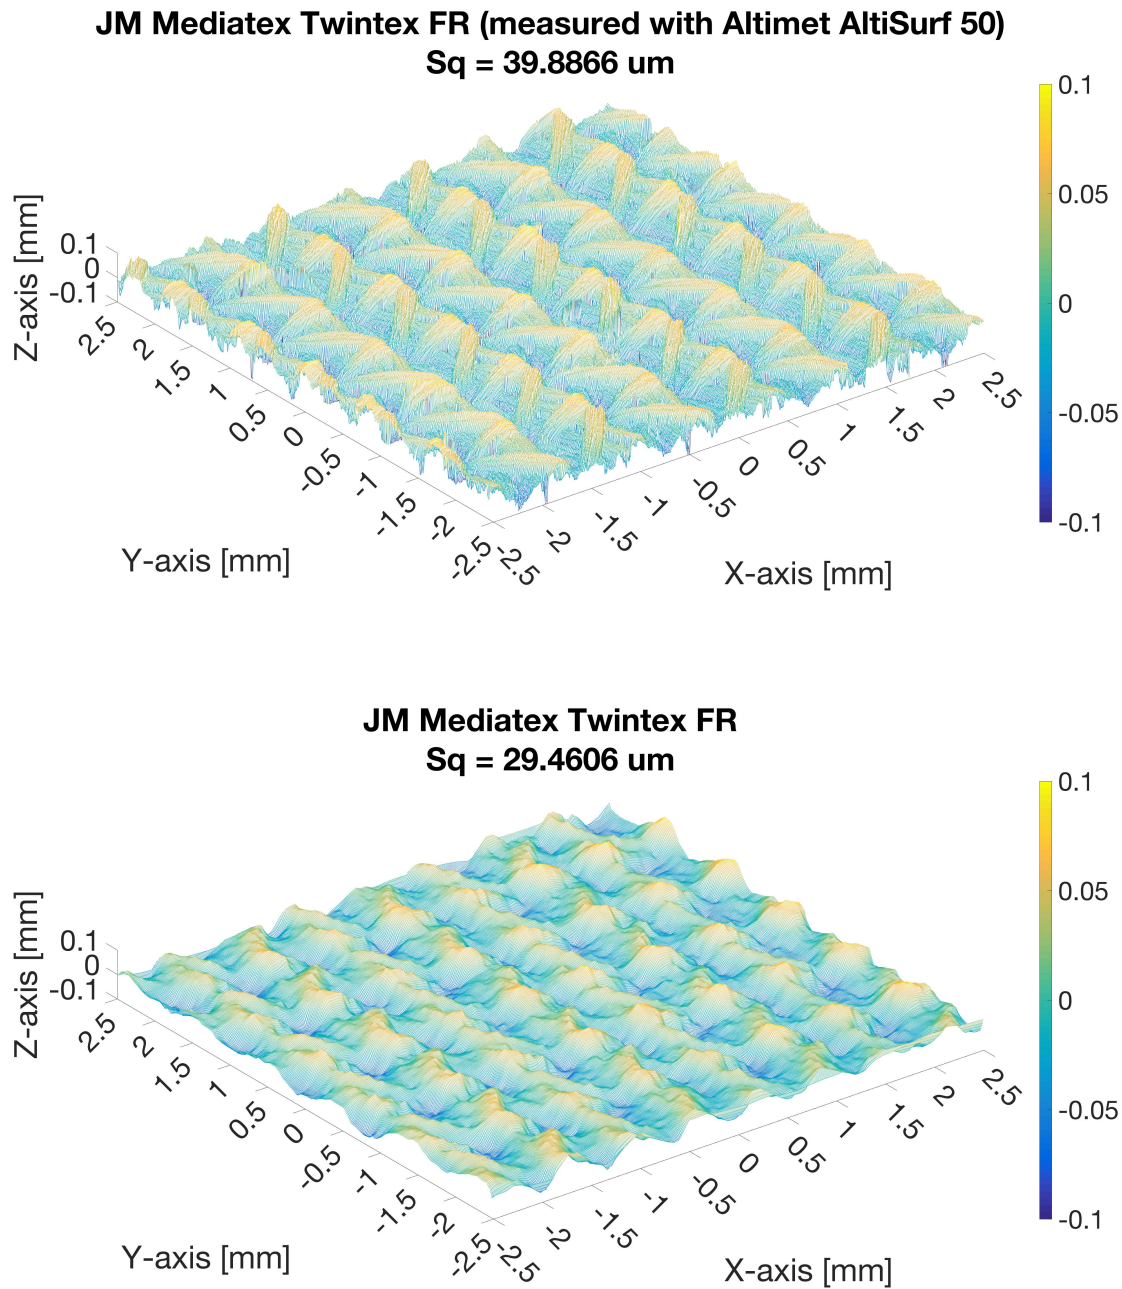


Figure 3.15 – Surface structure of JM Mediatex Twintex FR, polyester based knitted fabric. The colormap on the right side color codes the height of the surface points. The height is represented as the deviation from the best fitting plane to the scanned point cloud. All units are mm, except in the title, where they are μm . The top subfigure shows the point cloud computed by the Altimet AltiSurf 50, after we apply our simple compensation for the non-planarity of the sample. The bottom subfigure shows the point cloud computed with our structured light setup, after we apply the compensation for non-planarity of the sample and for the errors due to the radiometric calibration. Our scan shows lower surface roughness values since it cannot capture the high slopes and high-frequency details of the surface structure.

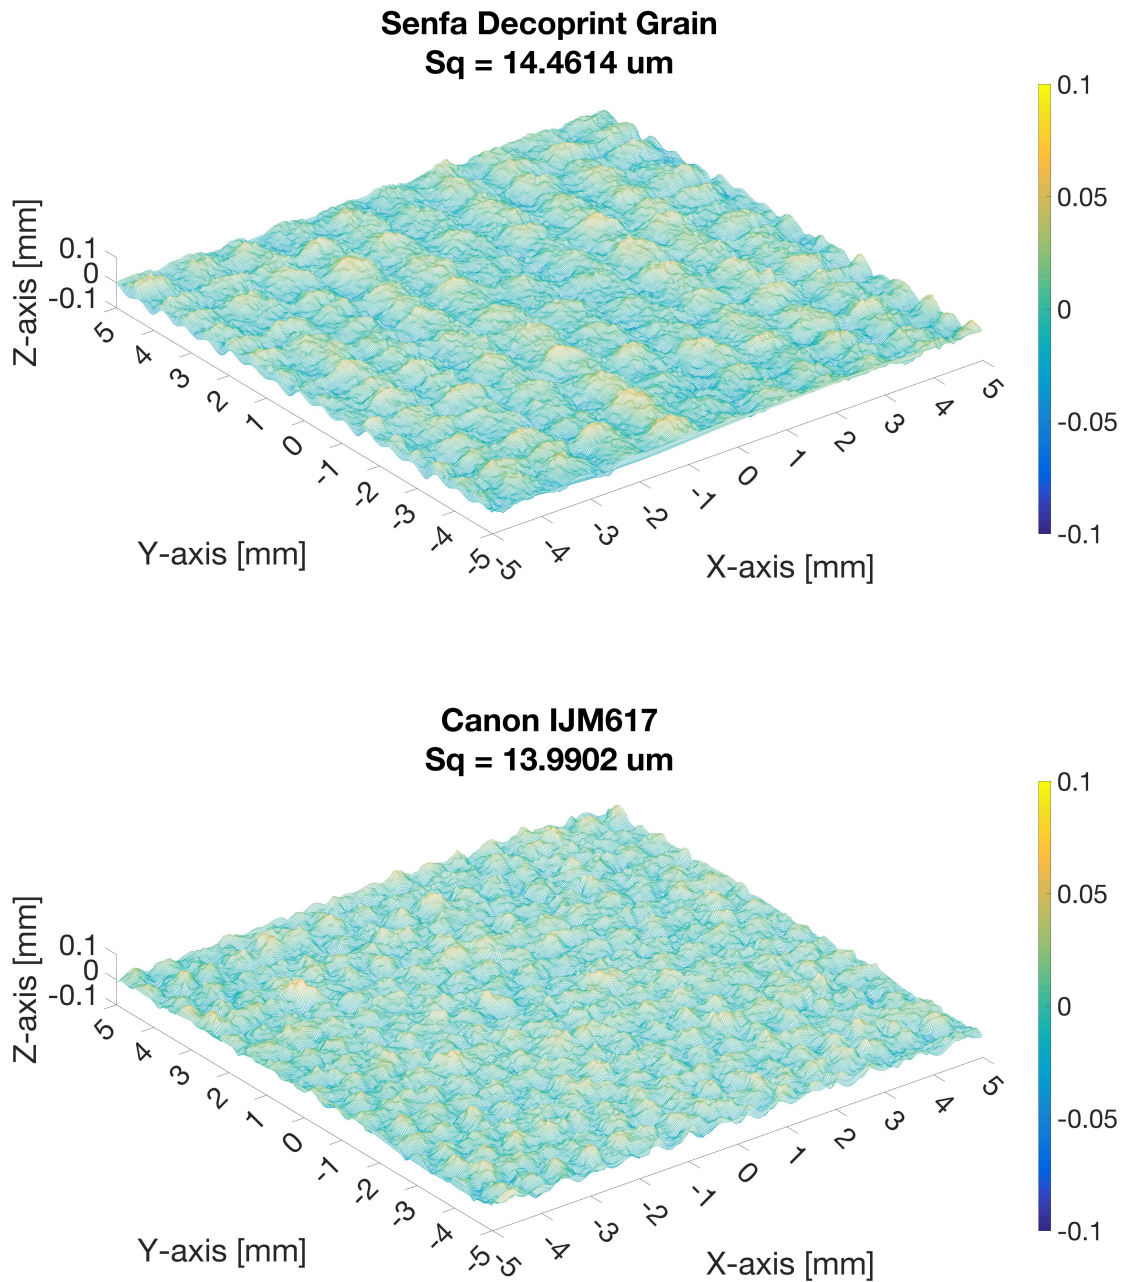


Figure 3.16 – Surface structure of Senfa Decoprint Grain (top), and Canon IJM 617 (bottom), for a 10 mm \times 10 mm scan. The colormap on the right side color codes the height of the surface points. The height is represented as the deviation from the best fitting plane to the scanned point cloud. All units are mm, except in the title, where they are μm . Both point clouds were computed with our structured light setup, after we apply the compensation for non-planarity of the sample and for the errors due to the radiometric calibration. The point clouds produce similar surface roughness values to the 5 mm \times 5 mm scans, although with slight increase due to global warping of the samples.




Figure 3.17 – Surface structure of Canon IJM 618 (top), and Senfa Decoprint Night (bottom), for a 10 mm \times 10 mm scan. The colormap on the right side color codes the height of the surface points. The height is represented as the deviation from the best fitting plane to the scanned point cloud. All units are mm, except in the title, where they are μm . Both point clouds were computed with our structured light setup, after we apply the compensation for non-planarity of the sample and for the errors due to the radiometric calibration. The point clouds produce similar surface roughness values to the 5 mm \times 5 mm scans, although with slight increase due to global warping of the samples.

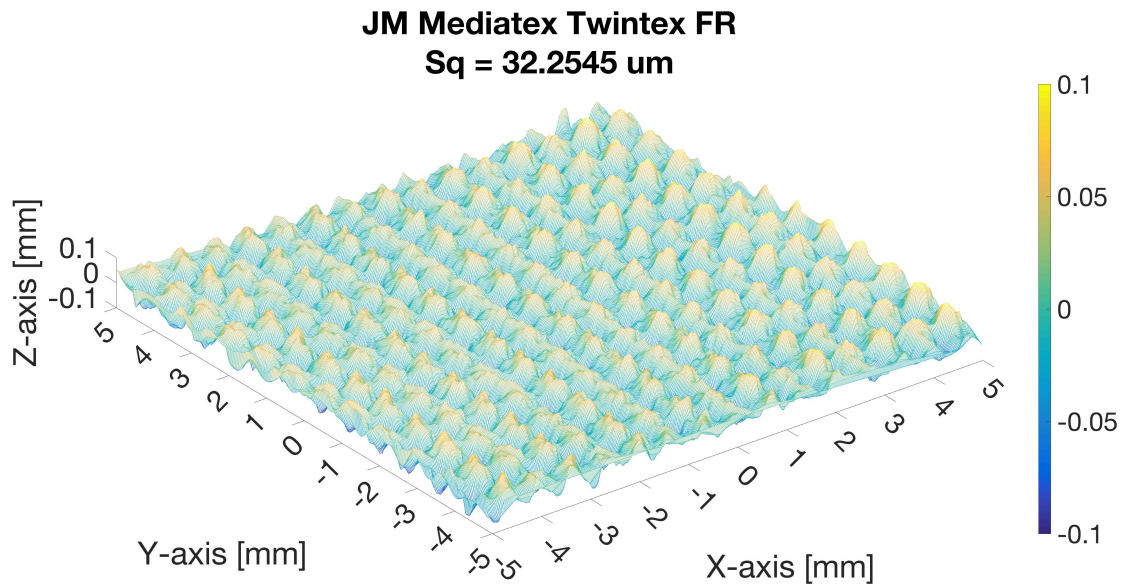


Figure 3.18 – Surface structure of JM Mediatex Twintex FR, for a 10 mm \times 10 mm scan. The colormap on the right side color codes the height of the surface points. The height is represented as the deviation from the best fitting plane to the scanned point cloud. All units are mm, except in the title, where they are μm . The point cloud was computed with our structured light setup, after we apply the compensation for non-planarity of the sample and for the errors due to the radiometric calibration. The point cloud produces similar surface roughness value to the 5 mm \times 5 mm scan, although with slight increase due to global warping of the samples. A non-linear warp of the surface can be seen in the right corner of the point cloud.

4 Comparative analysis of transmittance measurement geometries and apparatus

The accurate measurement of reflectance and transmittance properties of materials is essential in the printing and display industries in order to ensure precise color reproduction. In comparison with reflectance measurement, where the impact of different geometries ($0^\circ:45^\circ$, $d:8^\circ$) has been thoroughly investigated, there are few published articles related to transmittance measurement. In this work, we explore different measurement geometries for total transmittance, and show that the transmittance measurements are highly affected by the geometry used, since certain geometries can introduce a measurement bias.

In Section 4.1 we start by presenting relevant work that provides guidelines for measuring total transmittance. We then continue to discuss the different aspects of transmittance, and how they can be measured, in Section 4.2. We also explain formally why there are differences between the various measurement methods for total transmittance. In Section 4.3, we show our flexible custom setup that can simulate these geometries, which we evaluate both qualitatively and quantitatively over a set of samples with varied optical properties. In Section 4.4 we also compare our measurements against those of widely used commercial solutions, and show that significant differences exist over our test set. However, when the bias is correctly compensated, very low differences are observed. These findings therefore stress the importance of including the measurement geometry when reporting total transmittance.¹

4.1 Related work

Transmittance measurement is used in a wide range of industries, from solar cells [103] and light filter manufacturing [95] to quality control in food production [100]. The accurate characterization of transmittance is especially important for applications that depend on backlight illumination, which are created by depositing an image onto a substrate that is illuminated from the back side. Common backlit devices include LCD displays [54] and printed

¹The material in this chapter was presented at Electronic Imaging 2019 [91].




Figure 4.1 – Total transmittance of our diverse set of samples. The samples include smooth transparent plastics, milky plastics, papers and textiles. The measurements were performed with our custom setup in the $0^\circ:d$ geometry. Best viewed on color display.

advertisements [21]. The correct display on these devices requires accurate spectral prediction models, and transmittance measurements are used extensively in their calibration [43, 44, 59, 64].

Materials generally transmit light in both direct and diffuse manner, and both have to be captured for *total* transmittance measurement. The diffusion is caused by intrinsic light scattering and by rough material interfaces. This causes light to leave the material in all directions of the hemisphere, with angularly varying radiance.

The most common method for capturing *total* transmittance is by using integrating spheres, whose interior is coated with a highly reflective and almost Lambertian material. The light that leaves the sample after transmission has a non-uniform angular distribution, and therefore does not illuminate the sphere uniformly. However, after this direct illumination of the sphere, every sphere surface point re-illuminates the sphere due to reflection. Thus, irrelevant of its position, the detector captures a flux that corresponds to the transmittance of the sample, integrated over the hemisphere of directions. It is preferable to place the detector at a point where no direction path from the sample to the detector is possible, commonly by using baffles. Capturing total transmittance requires proper measurement setup and diligence. For that reason, several works and industrial standards have been developed that carefully outline the necessary steps for its capture [47, 27, 41, 50, 102]. All these approaches are based on integrating spheres.

In [103], Yu et al. compare the results of the above mentioned industrial standards for haze estimation, which depends on total transmittance measurement. Additionally, they compare their accuracy for measuring total transmittance. Apart from this work, the influence of measurement geometries remains too rarely addressed. We therefore present one of the first comparisons of different geometries for total transmittance measurement.

The main objective of our work is to quantify the impact that differences in measurement geometry have on the measured transmittance values. We are interested in measuring the total transmittance, i.e., both direct and diffuse components, as our goal is to measure the total amount of light that leaves the samples. We have therefore built a custom setup, which is easily modifiable and allows us to gain insights into how much the different effects, existing for different measurement geometries, affect the results. We evaluate their performance over a set of samples with different transmitting properties, whose total transmittances are shown in Figure 4.1. These samples include smooth transparent plastics, milky plastics, papers, and textiles. We show that the various geometries provide measurements with non-trivial differences.

Our secondary objective is to analyze which configuration yields results that are closest to those of commercial spectrophotometers, which employ various measurement geometries. The comparisons show that there are notable differences even between the commercial solutions. Some measurement geometries introduce a significant measurement bias, which should be compensated. Total transmittance thus depends on the measurement geometry and should be reported with it.

4.2 Measuring total transmittance

In order to define the transmitting properties, different aspects of material transmittance can be measured, such as *direct*, *diffuse*, and *total* transmittances, *haze* and *clarity*. Direct transmittance measures the amount of light that leaves the sample in the same direction as the incident beam direction, i.e., without being scattered. Diffuse transmittance measures the light that is scattered by the sample in all directions except the direct. Total transmittance is the sum of the direct and the diffuse components. Haze describes the wide-angle scattering (at an angle greater than 2.5° from incident beam direction [27]) and is calculated as the ratio of diffuse to total transmittance. Finally, clarity compares the light flux intensity subjected to narrow angle scattering (excluding 0°) to the directly transmitted flux.

Note that what we refer to as *transmittance*, might also be considered a *transmittance factor*. This is because we measure the transmitted light after a reflection from a Lambertian diffuser (the integrating sphere). However, the transmittance factor cannot be higher than 100 %, because the transmitted light is integrated over the hemisphere of directions for both the incident and transmitted fluxes. In the case of the reflectance factor, the values can be higher than 100 % when measuring specular materials, since the measurements are referenced to a Lambertian diffuser.

Chapter 4. Comparative analysis of transmittance measurement geometries and apparatus




Figure 4.2 – Comparative methods for measuring total transmittance maintain a constant sphere reflectance when taking the reference and sample measurement. (a) Double beam geometry. A second illuminating beam is used to perform the reference measurement. (b) Integrating sphere with compensation port. The sample is switched between the entrance and compensative ports for the sample and reference measurements, respectively. The detector not visible as it is located below the sphere in both geometries.

Direct transmittance can simply be measured by a collimated light source and a spectrometer. Attempting to measure total transmittance with an apparatus built for direct transmittance will capture only the direct component, and all scattered light will be lost. This will yield much lower transmittance values compared to the true sample transmittance for scattering samples. Therefore, the measurement of total transmittance (along with diffuse transmittance and haze) requires a more complicated apparatus, such as goniometer-based or integrating sphere-based methods.

Instruments capable of measuring total transmittance can employ *comparative* or *substitutive* methods. Within each method, there are different measurement geometries. Two representative geometries for the comparative method are shown in Figure 4.2. Figure 4.2a shows the schema of the double beam geometry. It uses a collimated or converging light beam to illuminate the sample that is placed at the entrance port of the integrating sphere. In addition, while the sample is still placed at the entrance port, there is a second unobstructed entrance through which an equivalent light beam is sent to take the reference measurement. This allows for the measurement of the sample and the reference while maintaining the same average sphere reflectance. Therefore, it produces highly accurate measurements.

Another measurement geometry within the comparative methods is shown in Figure 4.2b. It uses a single beam in combination with an integrating sphere that features a compensative port. The compensative port is empty while the sample measurement is taken, and it is populated by the sample when the reference measurement is taken, maintaining a constant average sphere reflectance.

The substitutive methods generally require simpler measurement setups, since they use a single beam and integrating spheres without additional ports, apart from the entrance and

measurement ports. However, this requires the samples to be removed from the entrance port when taking the reference measurement. Example geometries are featured in Figure 4.4. This substitution creates a difference in the average sphere reflectance between the two measurements, since the empty port has zero reflectance, while the sample has a non-zero reflectance. More concretely, the light flux that radiates from the sample gets reflected from the sphere wall to irradiate the back side of the sample, which then reflects back some of this light into the sphere. This light flux is read by the sensor and reported as a higher transmittance value. Corrections have been proposed [94], which require the hemispherical reflectance of the measured sample. They can be extended by incorporating the diffuse and direct components of the sample transmittance, and the degree of sphere wall diffusion [34].

4.2.1 Sphere radiance

The difference between these two classes of methods can be shown formally. The radiance L of an internally illuminated sphere has the following expression [6]:

$$L = \frac{\Phi}{\pi A} M, \quad (4.1)$$

where Φ is the incident radiant flux, A is the sphere surface area and π is the total projected solid angle from the sphere surface, assuming a Lambertian reflectance. If we multiply the fraction from this equation by a surface reflectance, we can compute the radiant flux of a diffusely reflecting surface. However, the radiance of the sphere is increased by a sphere multiplier M , because each sphere spot is reflective and therefore radiating. The sphere multiplier is computed as follows [103]:

$$M = \frac{\rho^o}{1 - \rho^s \left(1 - \sum_{i=0}^n f_i \right) - \sum_{i=0}^n \rho_i f_i}, \quad (4.2)$$

where ρ^o is the surface reflectance for the incident light flux, ρ^s is the reflectance of the internal sphere wall, ρ_i is the reflectance of port i , and f_i is the fraction of port i to the total sphere area. If the incident beam illuminates the sphere wall, $\rho^o = \rho^s$. Equation 4.2 accounts for the increase in radiance due to multiple reflections, and can also be regarded as the average number of bounces of a photon inside the sphere before it is absorbed or it escapes.

From Equation 4.2 we can see that the reference and sample measurements with the substitution methods are subjected to different average sphere reflectance. When taking the reference measurement, all sphere ports have a zero reflectance, thus the sum featuring ρ_i equals zero. However, when a sample is present at the entrance port, we must include its reflectance as the corresponding ρ_i . This yields a higher sphere multiplier M . In contrast, the comparative methods maintain the same average sphere reflectance for both reference and sample measurement.




Figure 4.3 – Our measurement setup consists of an illuminating optical fiber, an iris diaphragm or opal glass (configuration dependent), an integrating sphere and a capturing optical fiber connected to a spectrometer. Depicted is the configuration 0°:d.

Another observation from Equation 4.1 is that as we decrease the size of the sphere, A , we increase its radiance. However, this comes at a cost, since smaller spheres have a lower sphere multiplier due to their larger port fraction f . This makes it more difficult for smaller spheres to distribute (“integrate”) the radiance uniformly across their inner surface, and thus are more dependent on the incident light flux distribution.

4.3 Our measurement setup

Our approach for measuring the total transmittance follows the ASTM D 1003 standard [27]. We are restricted by the design of our integrating sphere to use a substitutive method. Nevertheless, our study compares the influence of different geometries given the same sphere design. We also compare the results from this modular setup to commercial measurement instruments (Agilent Cary 5000 and X-Rite Color i7).

For our experiments, we built a custom optical setup, whose main components are visualized in Figure 4.3. It consists of an illuminating optical fiber connected to a Dolan-Jenner DC-950-HR illuminator, an iris diaphragm or opal glass (configuration dependent), an integrating sphere and an optical fiber connected to a Maya2000 Pro spectrometer.

To approximately collimate the incident illumination, we use an iris diaphragm at a distance of 100 mm from the illuminating fiber, and 25 mm from the sphere. This creates a light beam of approximately 8 mm in diameter at the entrance port of the sphere. Due to subsurface scattering, the light can travel an additional 8.5 mm radially across the sample plane before overfilling the entrance port. To create a diffuse illumination instead, we exchange the iris diaphragm by a diffusing opal glass. This configuration illuminates the sample across its entire surface. Note that this results in an approximately diffuse illumination, since the opal glass' size is finite, and thus not all incident angles are present. The diffuse illumination in the third configuration (see Figure 4.4c) is achieved by both the opal glass and the integrating sphere, and it is therefore uniform. The illumination did not include UV radiation.

The integrating sphere has a diameter of 100 mm and includes one port of 25 mm and one port of 10 mm diameter. Thus, the port fraction of our sphere is $f = 1.84\%$. The two ports are located on its equator, separated by 90° . The sphere's inside is coated with a highly reflective and diffusing coating.

The output optical fiber has a 25.4° field-of-view ($\approx f/5$), and is protected from direct incident illumination by a baffle. The sphere does not have baffling against first bounce radiance. Therefore, we keep the fiber slightly outside of the integrating sphere, while keeping its field-of-view completely inside the sphere.

We connect the output fiber to the Maya2000 Pro spectrometer for capturing the data. We used a temporal averaging of 5 captures and we applied a spectral boxcar smoothing filter of 10 samples length. The spectrometer captures UV, visible and NIR information from 200 nm to 1120 nm, with a step of 0.5 nm, totaling 2068 samples. Since we are interested only in the visible part of the spectrum, we disregard the captured spectra outside of the range 400 nm to 700 nm, and lower the spectral resolution to 5 nm. This allows us to compare to the commercial spectrophotometers.

Before each set of measurements, we acquired a dark spectrum by blocking the sample port of the sphere with an opaque sample. We then alternatively captured a reference spectrum, and a sample spectrum. For the reference measurements we left the sample port empty. To get the transmittance of the measured sample, we subtracted the dark spectrum from both the reference and the sample spectra, and then divided the sample spectrum with the corresponding reference spectrum.

Table 4.1 – The measurement configurations (geometries) that we used to capture total transmittance with our setup.

Geometry	Illumination	Capture	Sample position
0°:d	collimated	diffuse	sphere entrance
d:d	diffuse	diffuse	sphere entrance
d:0°	diffuse	direct	sphere exit

Chapter 4. Comparative analysis of transmittance measurement geometries and apparatus




Figure 4.4 – The set of measurement geometries that we have evaluated with our custom setup. They enable us to capture total transmittance. The details of these geometries can be found in Table 4.1.

The three measurement configurations that we evaluated with our modular setup are listed in Table 4.1. Furthermore, they are graphically illustrated in Figure 4.4. In the $0^\circ:d$ geometry (see Figure 4.4a), we place the sample at the sphere entrance port (25 mm). It is then illuminated with a collimated beam of light. The sphere integrates the complete light flux that leaves the sample, and we read the radiance of the sphere wall with the spectrometer. In the $d:d$ geometry, shown in Figure 4.4b, we illuminate the sample with a partially diffuse incident light flux, which we create by placing the opal glass in front of the sample. Finally, in the $0^\circ:d$ geometry (Figure 4.4c), the sample is illuminated with a completely diffuse light flux, since we place it at the sphere's exit port (25 mm). Note that in $0^\circ:d$ we switch the entrance and exit ports to maintain equal area of the sample port across geometries.

In addition to studying the effects of different measurement configurations, we also compare our results to those from commercially available solutions, namely the Agilent Cary 5000 and the X-Rite Color i7. The Cary 5000 spectrophotometer uses a double beam and a $0^\circ:d$ measurement geometry. Our measurements were performed with its internal DRA (Diffuse Reflectance Accessory) that includes an integrating sphere of 110 mm diameter. The captured spectra contain information inside the 400 nm to 700 nm spectral range, with a resolution of 5 nm. Unlike our illumination source, the one in the Cary includes UV radiation.

The Color i7 uses a $d:0^\circ$ geometry. This geometry can also be used for measuring total transmittance due to reciprocal optical geometry, where the hemispherical detector's field-of-view ($0^\circ:d$) is replaced by a hemispherical input flux ($d:0^\circ$). This geometry has an advantage of greater incident light flux, since the integrating sphere collects most of the light from the illuminant. The integrating sphere has a diameter of 6 inch (152.4 mm). We again extract the measured information from 400 nm to 700 nm, and adjust the measurements' resolution from 10 nm to 5 nm by interpolation. The UV filter was inserted for all measurements.




Figure 4.5 – An example set of sample surface texture and appearance. Note that the test samples are very varied to allow us to properly evaluate the accuracy of the different measurement geometries and apparatus.

4.4 Experiments and results

For our experiments, we created a test set of 16 samples based on their varied reflecting, absorbing and scattering characteristics. The list of samples together with their total transmittance can be found in the legend of Figures 4.1 and 4.6, where they are sorted by their average transmittance. The samples are fairly devoid of color, since they are used as print media. In the following, we consider the first three samples to be transparent, and the rest to be diffusing.

An example set of the samples' surface texture can be seen in Figure 4.5. The IJM 683 is a glossy, transparent, adhesive plastic film. MPI 3151 is a satin, white, translucent, adhesive vinyl. The Lumitex 195 is a fine coated polyester fabric, and the Decomural Coating is a matt, textured vinyl with a very rough surface.

4.4.1 Inter-configuration differences

Table 4.2 compares the transmittance measurements of the set of samples obtained using our custom setup. The measurement geometries that we used are detailed in Table 4.1. We compare the differences over all samples, but we also show them for only the transparent and diffuse sets of samples. For computing the differences, we used the mean absolute error (MAE), root mean square error (RMSE), and ΔE_{00} error metrics. The transmittance values were expressed in percentage.

A more comprehensive comparison of the measurement differences can be found in Figure 4.6. It shows the difference for every sample, over the complete spectrum, when measured in the stated configurations. Positive values translate into higher transmittance measurements for the first measurement geometry. For example, in Figure 4.6a, positive values mean that the 0°:d geometry measured higher transmittance values than the d:d geometry, and the opposite is true for the negative values.




Figure 4.6 – Inter-configuration differences. These figures show the differences in transmittance between the different geometries of our transmittance measurement setup. Positive values translate into higher transmittance measurements for the first measurement geometry. The geometries are presented in Table 4.1. A numerical breakdown of the accuracy can be found in Table 4.2. Best viewed on color display.




Figure 4.7 – Inter-spectrometer differences. These figures show the difference in transmittance between the commercial spectrophotometers and our setup in the $0^\circ:d$ geometry. Positive values translate into higher transmittance measurements for the first measurement apparatus. A numerical breakdown of the accuracy can be found in Table 4.3. Best viewed on color display.

We can see that the difference between the configurations that use diffuse capture is not very significant, although we change the illumination from collimated to diffuse. The differences are slightly more pronounced for the diffuse samples than for the transparent samples. In Figure 4.6a we can see that there is a bias in the measurements. The d:d geometry consistently measures higher transmittance for the diffusing samples, with a median difference of -0.84 %. In the remainder of the text, we use the term *bias* to refer to the median difference across a set of measured samples. If we remove this bias, i.e., subtract 0.84 % from all d:d measurements, the MAE decreases to 0.85 %. In these two configurations, the average sphere reflectance is consistent between the measurements of the same sample, and the only difference is in the distribution of the incident light flux.

In the second row of Table 4.2 we notice that the MAE between 0°:d and d:0° is less than 1 %. On a closer inspection, we see that the transparent samples show larger differences than the diffuse samples. Figure 4.6b shows us that two of the three transparent samples appear to have higher transmittance when measured with d:0°. On the other hand, we note that the spread of errors for the diffuse samples is low, thus the measurements are consistent. The bias across all samples is 0.06 %. If we remove MPI 2040, IJM 778, and IJM 617, we arrive to a MAE of 0.48 %.

It should be noted that when measuring transparent samples, the effective geometry for both 0°:d and d:0° is 0°:0°. This is because for 0°:d, the radiance that the sample transmits has the direction of the incident beam, i.e., it is at 0°. For d:0°, the reverse is true; the detector captures the radiance that leaves the sample at 0°, which is the result of the sphere radiance that is an extension of the detector's field-of-view. The collimation of light might also add to the differences.

The differences between d:d and d:0° are the most pronounced. In Figure 4.6c, we see both previous effects in this case, i.e., the transmittance of the transparent samples is greater for d:0°, and there is a measurement bias of 0.85 % for d:d. These factors contribute to the greatest MAE of 1.31 %. After correcting for the bias, the MAE decreases to 1.02 %.

All geometries maintain equal average sphere reflectance when measuring the same sample. From this we conclude that the measurement bias is inherent to the d:d geometry. It is likely caused by multiple reflections between the samples and the opal glass. These reflections do

Table 4.2 – Inter-configuration differences. This table compares the measurements performed with the geometries from Table 4.1. The results are evaluated over the set of samples shown in Figure 4.1.

Geometry	All samples			Transparent			Diffuse		
	ΔE_{00}	RMSE	MAE	ΔE_{00}	RMSE	MAE	ΔE_{00}	RMSE	MAE
0°:d vs. d:d	1.05	1.35	1.11	0.32	1.06	0.94	1.22	1.41	1.15
0°:d vs. d:0°	1.24	1.24	0.81	0.49	1.99	1.56	1.41	0.99	0.64
d:d vs. d:0°	1.44	1.61	1.31	0.66	2.06	1.77	1.62	1.48	1.20

Chapter 4. Comparative analysis of transmittance measurement geometries and apparatus

not occur in the other two geometries (see Figure 4.4). Therefore, care must be taken when positioning the elements in the measurement setup.

4.4.2 Inter-spectrometer differences

Table 4.3 summarizes our comparisons to commercial spectrophotometers. We have selected the 0°:d geometry of our setup for these comparisons, because it showed the best accordance to both the Cary 5000 and the Color i7. Here it is interesting to observe that the commercially available spectrophotometers have a similar measurement difference between their measurements as well as to those from our setup.

In Figure 4.7a, we can see that there is a negligible bias between the Cary 5000 and the Color i7 measurements, amounting to 0.03 %. An obvious error source can be seen in the shorter wavelengths, between 400 nm and 450 nm. This error stems from the fact that the measurements with the Cary 5000 were performed under illumination that included UV radiation, however, for the Color i7's measurements the UV radiation was filtered out. These differences come from fluorescence, since some of the samples fluoresce. If, for a better comparison, we consider only the wavelengths from 500 nm to 700 nm, then the MAE comes down to 1 %. By removing the bias, we see an insignificant MAE decrease to 0.99 %.

Continuing to Figure 4.7b, we can see that there is a prominent measurement bias. Since our 0°:d setup uses a substitution measurement method, this bias is caused by the non-constant average reflectance of the integrating sphere between the reference and sample measurement. The light that reflects from the sample's interface back into the sphere increases the measured transmittance, since we measure the radiance of the sphere wall. The bias amounts to -1.95 %, i.e., 1.95 % higher median transmittance for 0°:d. When we remove it from the measurements, the MAE becomes 1.58 %, which is comparable to the Cary 5000 vs. Color i7 difference. In this comparison, we again have the errors in the short wavelengths caused by the UV illumination of the Cary 5000. If we compare the measurements in the 500 nm to 700 nm wavelength range, we obtain a MAE of 2.22 %, or 1.07 % after also removing the bias of the 0°:d measurements.

Finally, in Figure 4.7c we compare the measurements from the Color i7 and our 0°:d geometry. We again see the effects already encountered in the previous comparisons. Namely, the 0°:d has a measurement bias of 1.75 %, somewhat lower when compared to that of the Cary 5000.

Table 4.3 – Inter-spectrometer differences. This table compares the measurements performed with two commercial spectrophotometers and our measurement setup in the 0°:d geometry. The results are evaluated over the set of samples shown in Figure 4.1.

Spectrometer	All samples			Transparent			Diffuse		
	ΔE_{00}	RMSE	MAE	ΔE_{00}	RMSE	MAE	ΔE_{00}	RMSE	MAE
Cary 5000 vs. Color i7	1.76	2.92	1.48	0.38	1.27	0.83	2.07	3.18	1.63
Cary 5000 vs. 0°:d	2.75	3.51	2.62	0.72	2.65	2.30	3.22	3.68	2.69
Color i7 vs. 0°:d	1.82	2.09	1.85	0.67	2.77	2.62	2.08	1.90	1.68

Table 4.4 – Updated inter-spectrometer differences. This table updates the results from Table 4.3 by considering the various sources of error in order to provide a better comparison between the measurement apparatus. The results are evaluated over all samples. The effects shown in the table header are cumulative.

Spectrometer	Original		500-700nm		Bias	
	RMSE	MAE	RMSE	MAE	RMSE	MAE
Cary 5000 vs. Color i7	2.92	1.48	1.33	1.00	1.31	0.99
Cary 5000 vs. 0°:d	3.51	2.62	2.45	2.22	1.33	1.07
Color i7 vs. 0°:d	2.09	1.85	2.01	1.77	0.96	0.76

However, we do not see the errors in the short wavelengths, since both measurements were performed in the absence of UV radiation. If we correct for the measurement bias, we achieve a MAE of 0.79 %. For a complete comparison, we also compute the MAE over the 500 nm to 700 nm range, which amounts to 0.76 %.

An updated comparison table that takes into account the previously discussed effects is shown in Table 4.4. The effects listed in the table header are cumulative. They represent the original results over all 16 test samples, evaluating the differences on the 500 nm to 700 nm wavelength range, and removing the measurement bias. The final column (Bias) incorporates the effects from the previous one. Although the start and end differences are comparable, the effects that exist for the various measurement setups are different. The table outlines the bias that exists in our implementation of the 0°:d geometry, which is based on a substitutive method.

4.4.3 Repeatability

We compared the measurements of the samples done with our 0°:d setup over a span of more than one month, and note that there are negligible differences, comparable to those of the commercial spectrophotometers. We report the differences in Table 4.5. The MAE between our measurements in the 0°:d configuration performed on 29.03.2018 and 09.05.2018 is 0.33 %. The MAE between the measurements of the same set of samples done with the Color i7 spectrophotometer done on 29.11.2017 and 14.03.2018 is also 0.33 %.

Table 4.5 – Measurement repeatability. We report the average difference over the same set of measured samples that were taken with a time difference of at least one month.

Spectrometer	ΔE_{00}	RMSE	MAE
0°:d	0.58	0.48	0.33
Color i7	0.37	0.39	0.33

4.5 Summary

In this work, we discussed the effects that influence the accuracy of total transmittance measurements for various measurement methods and geometries that use integrating spheres. To show the differences in practice, we constructed a custom optical setup that can be configured to the $0^\circ:d$, $d:d$, and $d:0^\circ$ geometries. These were based on a single beam substitutive measurement method. We evaluated their accuracy over a test set of 16 samples that have various levels of transmittance, absorbance and reflectance. The measurements showed a MAE as low as 0.81 % between $0^\circ:d$ and $d:0^\circ$, and up to 1.31 % between the $d:d$ and $d:0^\circ$ geometries. We explained why the effective geometry for transparent samples for $0^\circ:d$ and $d:0^\circ$ is actually $0^\circ:0^\circ$, whereas for the $d:d$ it is indeed $d:d$. We noted that the $d:d$ geometry consistently measured higher transmittances than $0^\circ:d$ and $d:0^\circ$, whose measurements did not show an apparent bias. This is likely caused by our implementation of the geometry, which introduced multiple reflections between the sample and the opal glass. Therefore, care must be exercised when building custom measurement setups.

We also compared the results obtained by our setup to those of commercial spectrophotometers, the Agilent Cary 5000 and the X-Rite Color i7. The $0^\circ:d$ geometry had the best measurement accordance to the commercial solutions, which is also the geometry that the Cary 5000 uses. However, the Cary has a double-beam configuration (comparative method), and our setup has a single beam configuration (substitutive method). At first glance, all spectrometers showed similar differences in their measurements. However, those inconsistencies were caused by different effects. Our $0^\circ:d$ geometry showed a measurement bias because of varying average reflectance of the integrating sphere during reference and sample measurements. This is caused by the nature of the substitution methods, and in our experiments lead to a MAE of up to 2.22 %. After correcting for these effects, all apparatus performed to a MAE of 1.07 % or less.

When reporting total transmittance it is thus good practice to also include the measurement geometry. As we have shown, some measurement geometries introduce a significant bias. However, if it is compensated, very low errors can be achieved.

5 BTDF Estimation with Radial Imaging


The Bidirectional Scattering Distributional Function (BSDF) describes how light is distributed after an interaction with a material. It is made up of two components, the Bidirectional Reflectance Distribution Function (BRDF) and Bidirectional Transmittance Distribution Function (BTDF), which, in turn, describe the scattering and the transmission of light, respectively.

The BSDF is of great importance in many domains. Some examples include rendering, visual special effects, computer games, product design, advertising and cultural heritage, where one is interested in perfecting the appearance of objects. It can also have a functional application in modelling of stray light in optical systems and in architectural modelling of internal illumination of spaces.

The BSDF has traditionally been measured with goniophotometers [45]. They generally consist of an arm capable of rotating around two axes in fine angular increments. A sensor is attached to the end of the arm that measures the reflected, transmitted and/or emitted light by the measured sample, which stays fixed. This process captures the respective quantities for the desired outgoing directions, and the process is then repeated after the sample is illuminated from a different direction. An example modern goniophotometer, the Pab Advanced Technologies Ltd pgII, can be seen in Figure 5.1a.

Although the scanning goniophotometers are very accurate, their scanning times can be quite long, due to the mechanical nature of the process. To decrease the scanning time, other alternatives have also been proposed, often called parallel goniophotometers [51]. These devices can simultaneously detect many, or even all scattering directions. To achieve this parallelism, they use dioptric elements (lenses) or catadioptric elements (diffuse or mirrored reflectors) to manipulate the scattered light, and send it towards a camera that captures the angular distribution of the incoming light. An example parallel goniophotometer can be seen in Figure 5.1b.

The goal of this work is to measure a transmissive sample's BTDF. More specifically, our goal is to evaluate the feasibility of a compact, affordable and easy to build alternative to the established measurement techniques. We would like to study whether such a device is viable, and do not expect to reach the same level of accuracy. In addition, our proposed system



(a) The scanning goniophotometer pgII developed by pab Ltd, based on the work of Apian-Bennewitz [5]¹ (b) A parallel goniophotometer developed by Ghosh et al. [28]

Figure 5.1 – An example (a) scanning and (b) parallel goniophotometer. The scanning goniophotometers are based on a mechanical movement of the light source, sample and/or sensor to cover the directions of the incoming and outgoing hemispheres. The parallel goniophotometers can capture many outgoing directions in a single photograph, which can lead to a substantial speed up in the BSDF acquisition.

is flexible, and can be scaled from a small and portable device, to a large device capable of measuring larger samples, as needed in architectural light modelling [2].

A schematic diagram of the proposed device is shown in Figure 5.2. The idea is to use a cylinder made out of Spectralon, because it is the most Lambertian material known. The cylinder has the transmitting sample next to one of its openings, and a photo camera next to the other. The sample is illuminated by a collimated beam of light, the cylinder collects the transmitted light, and the camera captures the equilibrium radiance reflected by the internal cylinder wall.

Our approach is based on simulating the *infinite* light bounces inside the cylinder with a light transport matrix. By inverting this matrix we can compute the light that the camera would capture after a single light bounce, i.e., the light that leaves the sample towards the cylinder and bounces towards the camera. From the first bounce image, we can infer the light distribution of the sample in all outgoing directions that the cylinder covers. Although the computation of the light transport matrix is expensive, this process takes place offline and does not change as long as the geometry of the setup does not change. Therefore, once computed, it is applied through a simple matrix-vector product with the camera captured radiance, from where we can easily compute the BTDF of the measured sample.

We apply our approach to synthetic images generated by the physically based renderer Mitsuba [49]. With Mitsuba we have a complete control over the scene parameters and over the number of light bounces inside the cylinder, i.e., it allows us to set the BTDF of the sample, the BRDF of the cylinder coating, and the parameters of the camera.

¹Image taken from: <https://rgl.epfl.ch/pages/lab/pgII>




Figure 5.2 – Our proposed design for measuring a sample's BTDF: The elements on the left side of the schema, before the sample, are used to collimate the incoming light from the light source. The sample is illuminated by the collimated beam of light, and the transmitted light is collected by the diffusing cylinder. The steady-state radiance of the cylinder inside, after many light interreflections, is measured by the camera. By inverting the light transport inside the cylinder we are able to retrieve the BTDF of the sample.




Figure 5.3 – An isometric view of our proposed design for measuring a sample's BTDF: The elements have been cut in half for clarity.

When we assign a diffuse (Lambertian) BRDF to the cylinder, we know that the outgoing radiance from each facet towards the camera is the same as in any other direction. This allows for a very simple formulation and solution to the problem. Spectralon deviates substantially from being a Lambertian reflector at steeper incident angles, and makes the previous solution invalid. Therefore, we formulate and compute the more involved solution, that can be applied to real cylinders.

We built a prototype of the schema, shown in Figure 5.2, to test the performance of the setup in a real-world non-ideal conditions. Figure 5.3 shows a render of the proposed design. The accuracy of the system was severely degraded by alignment issues and problems with back-scattered light. However, we believe that a industrial grade implementation of the setup can overcome most of these issues, and provide the accuracy needed for certain less demanding applications, for example, profiling the angular distribution of light of monitors or the appearance of backlit prints and advertisements.

In Section 5.1 we provide an overview of the classes of scanning and parallel goniophotometers. We then introduce the basic radiometric quantities, which we will need for the later discussions, in Section 5.2. In Section 5.3 we derive the light transport equation (LTE), its solution operator, we show its surface form, and an analytical solution to the LTE for a simple scene with Lambertian materials. We then introduce the light transport matrix (LTM) in Section 5.4, which encapsulates the light transport for complex scenes. The LTM allows us to describe the light transport inside our cylinder for arbitrary materials. We then use it to reverse this complex light transport, and compute the BTDF of the sample.

We then continue to evaluate our proposed approach. Section 5.5 explains the algorithmic steps. In Section 5.6 we recreate our proposed goniophotometer design in the physically based renderer Mitsuba, and apply the algorithm on rendered image. We then build a prototype of the design using off-the-shelf optical component, and evaluate it in Section 5.7. We talk about some of the the practical considerations and limitations of the proposed design in Section 5.8. Finally, we conclude with a summary in Section 5.9.

The measurement configurations that we evaluated in this work are shown in Table 5.1. Each cell of the table links to the subsection where the configuration was evaluated.

Table 5.1 – This table summarizes the measurement configurations that we evaluated in this work. The main header represents the maximal number of bounces that the light was allowed to make inside the cylinder. The first column lists the cylinder coating types, and the subheader tells whether the sample was synthetic or measured. The cells of the table show the subsections where the corresponding configuration was evaluated.

Sample	Single light bounce		Infinite light bounces	
	Synthetic	Measured	Synthetic	Measured
Lambertian cylinder	5.6.3	5.6.3	5.6.4	5.6.4
Spectralon cylinder	5.6.3	5.6.3	5.6.4	5.6.4 / 5.7.3

5.1 Related Work

Due to its wide range of applications, many devices for capturing the BSDF have been designed and built, that use different measurement approaches. Some devices have a wide field of application, while others are purpose built for an application. They also differ in terms of size, accuracy, acquisition time, price, size of sample, ease of calibration, etc. In this section we will discuss the general concepts behind the major classes of devices. For much more thorough summaries of the different techniques, and the differences within the classes, we refer the reader to [2, 36, 51].

As outlined by Dupuy and Jakob [22], the main challenges for BSDF acquisition are its high dimensionality, arbitrary frequency and the physical constraints imposed on the equipment. The BSDF is considered to be a five-dimensional function, defined in terms of an incident and outgoing directions, each represented with two angles, and a wavelength. Acquiring such a high-dimensional data is costly due to the curse of dimensionality. Additionally, the BSDF can have an arbitrary frequency, which can cause abrupt changes in both the directional and spectral domains. This is difficult to know without dense measurements. Finally, physical phenomena such as stretching of light at grazing angles, beam occlusion, high dynamic range, etc., make the BSDF acquisition with physical devices even more challenging.

One of the earliest proposed goniophotometers by Hsia and Richmond [45] is composed of an articulated sample holder placed on a turntable, and a rotating arm that houses the detector. An example of a modern device, the Pab Advanced Technologies Ltd pgII, is shown in Figure 5.1a [5]. This device has an immovable light source, an articulated sample holder and a rotating arm, that can both rotate around 2 axes. The rotating arm houses photo diodes, a spectrometer, or a camera as a capture device. This setup covers 4 degrees of freedom of the BSDF, two incoming and two outgoing angles, and the wavelength can also be included if a spectrometer is used as a detector. The measurements are performed with a scanning process, where the capture arm rotates around the sample, while the sample is stationary. Then the sample is moved to the next position to be illuminated from a different incident direction, and the capture arm again covers all outgoing directions.

The scanning goniophotometers are the most accurate and reliable devices for capturing the BSDF. They, however, require long acquisition times because they capture a single measurement at a given position, and for repositioning they require a physical movement of the capture or illumination arms. Algorithms that optimize the capture times, by optimizing the arm trajectory, have been developed. Apian-Bennewitz [5] first introduced adaptive refinements in angular resolution. This has the task of assigning more samples around peaks and less in smooth areas of the BSDF. However, it requires a two-step process, where the first scan is used to identify the locations of such peaks, which in-turn increases the required time for scanning. A recent work by Dupuy and Jakob [22] also aims at optimizing the trajectory of the capture arm. It is also a two-step process, where the first step is quite efficient, as it uses the retroreflection of the measured material. Although this algorithm greatly reduces the




Figure 5.4 – Schematic diagram of the three main working principles of parallel goniophotometers. (a) Screen imaging devices use a backscattering screen to collect all outgoing light from the sample, and reflect it back in all helispherical directions, including towards the camera. (b) Dioptric angular mapping uses complex lens assemblies to direct all outgoing light from the sample towards a camera. (c) Catadioptric angular mapping uses an elliptical mirror to reflect all outgoing light from the sample, which is placed at one focal point, towards a camera, placed in the second focal point. Figure inspired by [51].

acquisition time, a high-quality scan with a good angular resolution still requires 2.5 hours for isotropic and 2-3 days for anisotropic samples, according to the authors.

Parallel goniophotometers have therefore been developed to decrease the measurement times. These devices can generally be classified as image based, dioptric or catadioptric [51]. To parallelize the capture process, these devices use a camera. This allows them to measure many outgoing directions at once, up to the full hemisphere.

The image based devices either capture the sample directly, or a projection screen instead. When directly capturing the sample, the camera has to be able to see multiple outgoing directions, which is why the sample has to have a non-planar and known shape, e.g., cylindrical or spherical [61, 63]. The indirect approaches capture a projection screen made out of a diffusing material. The screen can have a hemispherical [66], box [20] or flat [1, 69]. An example screen imaging setup, similar to the one proposed by McNeil and Wilson [66], can be seen in Figure 5.4a. The screen's purpose is to backscatter the outgoing directions from the sample towards the camera, which is why it is usually coated with a highly diffuse material. Due to the concave shape of the screens for most methods, interreflections occur, and they have to be compensated for. Furthermore, to measure the BRDF, a hole in the screen might be necessary for the incident beam to pass through, which leaves some outgoing angles unavailable. These systems can also have lengthy calibration procedures.

The dioptric devices use only light refraction to guide the outgoing sample directions towards the camera, i.e., to perform the angular (outgoing directions) to spatial (camera sensor) mapping. They rely on an optical Fourier transform performed by a simple lens, when the sample and the sensor are at its first and second focal planes, respectively. The intensities on the second focal plane are the coefficients of the Fourier transform of the spatial light distribution on the first focal plane, which relates to the outgoing directions [32]. To capture a wide range of outgoing directions, however, complex compound lenses that surround the sample have to be used. This approach is most commonly used to characterize the optical properties of liquid crystal displays. Devices based on this method can achieve high angular resolution and accuracy, but are expensive, and hard to develop and require a new design for each wavelength. An example dioptric parallel goniophotometer, can be seen in Figure 5.4b.

Catadioptric devices use polished hemispherical [98, 3], parabolic [17, 81] or elliptical [82, 40, 62, 92] mirrors to reflect many outgoing light directions towards the camera. An example catadioptric parallel goniophotometer, can be seen in Figure 5.4c. The sample is illuminated from either a hole in the mirror, through the mirror if it is semi-transparent, or with a beam-splitter for the parabolic design. In the elliptical mirrors, which is the most popular design, the sample is placed at one focal point, and the camera is placed at the other focal point. The light scattered by the sample is then reflected by the mirror towards the camera. These devices can provide fast scanning and do not have problems with stray light. One of the main downsides is that for larger acceptance angles ($> 40^\circ$) they require fish-eye lenses, which has been shown by Karamata and Andersen [52, 53] that the distortions and aberrations due to the lens make the systems severely biased.

Spherical and hemispherical gantries have also been used to recover the BSDF. The most prevalent design is a hemisphere filled with LEDs, a camera at the apex of the hemisphere, and a sample placed at the center of the hemisphere [60, 29]. Variations have also been proposed where the LEDs are used for both illumination and measurement [10], or the LEDs are replaced with cameras, whose flashes are used for illumination [83]. These devices were also used to estimate the specular and diffuse albedo, index of refraction, 3D geometry, to which they fit a BSDF model that can vary across the image, thus Spatially Varying BRDFs (SVBRDFs) [30, 37, 96].

Kuthirummal and Nayar [55] used a mirrored cylinder to capture multiple views of an object with a single camera. By setting the camera close to the cylinder and using a wide angle lens, they were able to capture multiple reflections of a sample. The captured data for the known viewing (outgoing) directions together with the incident direction of the light source is then used for fitting an analytical BRDF model.

Our approach draws inspiration from the radial imaging paper by Kuthirummal and Nayar [55]. It can also be classified in the category of imaging methods, closest to idea presented in [69]. We rely on an imaging screen, a cylinder made out of Spectralon, which is the most diffusing material known. While [55] used a polished cylinder and could capture the BRDF of

samples at discrete angles, the diffuse cylinder allows us to capture many outgoing directions in a single acquisition. We do not capture the complete hemisphere of outgoing directions, because of the cylinder opening at the camera end (Figure 5.2).

5.2 Radiometry

Radiometry deals with the measurement of electromagnetic radiation in the visible, ultraviolet and infrared wavelengths. Here we will define the quantities that we will need in the following sections. We start with the geometric quantities used in light transport. We then present the radiometric quantities. For a more detailed explanation of all these quantities, please see [65, 79].

5.2.1 Geometric quantities

We will use ω to denote a direction in \mathbb{R}^3 . We assume that the directions are unit vectors and that they are expressed relative to the surface normal at point p . We will denote the set of all directions with \mathcal{S}^2 .

Solid angle is the generalization of a planar angle to 3 dimensions. It is defined as the surface area that a given object subtends from the unit sphere centered at a point p . Alternatively, it can be viewed at the set of directions that the object occupies from the unit sphere centered at point p :

$$d\omega_p = \frac{\cos\theta \, dA_i}{r^2} \quad [\text{sr}], \quad (5.1)$$

where θ is the polar angle between the normal of the object and the vector joining its center and point p , and r is the length of the vector joining point p and the center of the object. The $\cos\theta$ term accounts for the decrease in the object's surface area dA_i when viewed from point p , due to the object's relative rotation, and the r^2 term projects the object to the unit sphere, from the sphere with radius r . The solid angle is expressed in the unit steradian (sr), and it is dimensionless.

Projected solid angle is computed by projecting the solid angle on the surface perpendicular to the normal at p [73]:

$$d\omega_p^\perp = |\omega \cdot \mathbf{n}_p| \, d\omega = \cos\theta_i \, d\omega = \cos\theta_i \frac{\cos\theta \, dA_i}{r^2} \quad [\text{sr}] \quad (5.2)$$

where θ_i is the angle between ω and \mathbf{n}_p , i.e., the surface normal at p .




Figure 5.5 – The Nusselt analog shows the steps of projecting the solid angle visually².

The steps for projecting the solid angle are illustrated in Figure 5.5. The projection onto the unit hemisphere takes care of $\cos\theta$ and r^2 , and produces the solid angle subtended by the object's surface area. The projection onto the disk centered at p gives the projected solid angle.

5.2.2 Basic quantities

Here we introduce some of the most important radiometric quantities for light transport, and also rendering, that we will use in the later sections: energy, flux, irradiance/radiant exitance and radiance. All of these quantities can be derived from energy by derivation, or from radiance by integration. The radiometric quantities are generally wavelength dependent, but we will omit this parameter for simpler expressions.

Energy that the photons of the electromagnetic radiation of interest carry is expressed as:

$$Q = \frac{hc}{\lambda} \quad [\text{J}] \quad (5.3)$$

where $h = 6.626 \times 10^{-34}$ J/s is Planck's constant, $c = 299,792,458$ m/s is the speed of light in vacuum, and λ is the wavelength of the electromagnetic radiation of interest. It is measured in joules (J).

Flux (or power) measures the amount of energy that passes through a surface per unit time:

$$\Phi = \frac{dQ}{dt} \quad [\text{W}] \quad (5.4)$$

²Image by Jheald: <https://commons.wikimedia.org/w/index.php?curid=15859738>.




Figure 5.6 – Basic radiometric quantities. (a) The flux that crosses the two spheres is equal, but the irradiance is different due to the difference in the surface area of the two spheres. (b) Compared to irradiance, radiance discerns between directions of light travel. Images taken from [79].

It is measured in watts (W).

Figure 5.6a shows the flux from a point light source that crosses two spheres. The flux measured on crossing either sphere is the same because it measures the total amount of energy. Alternatively, if we put two flat surfaces with equal areas and at a different distance from the point light source, the surface closer to the point light source will register a higher flux than the one further away, since the energy is more dense when it crosses the first surface. This leads us to the next quantity.

Irradiance is the total flux that falls onto a surface from all directions, divided by that surface's area A , i.e., it is the average density of power incident on A :

$$E(p) = \frac{d\Phi(p)}{dA} \quad \left[\frac{W}{m^2} \right] \quad (5.5)$$

If we go back to Figure 5.6a, the larger sphere measures a lower irradiance than the smaller sphere. This is because the point light source emits equally in all directions, and as the second sphere has a greater surface area than the first sphere, the emitted energy spreads over a greater surface, thus its density decreases. Similarly, in the example with the two flat surfaces, the one further away from the point light source will measure a lower irradiance.

Radiant Exitance, M , is defined equivalently to irradiance, E , however, it is used to denote the total flux *leaving* a surface in all directions. The measurement unit is the same as for irradiance, watts per meter squared (W / m^2).

Radiance is the most interesting quantity for light transport calculations. It adds one more dimension to the irradiance. It denotes the flux that falls on or leaves a surface, but it also discerns between directions (solid angles) in which the light travels. It is defined as:

$$L(\mathbf{p}, \omega) = \frac{dE_\omega(\mathbf{p})}{d\omega}, \quad E(\mathbf{p}) = \int_{\Omega} L(\mathbf{p}, \omega) d\omega \quad \left[\frac{\text{W}}{\text{sr m}^2} \right], \quad (5.6)$$

where E_ω is the irradiance at the surface perpendicular to the direction ω . The radiance can also be expressed in terms of flux:

$$L(\mathbf{p}, \omega) = \frac{d^2\Phi}{d\omega dA^\perp} = \frac{d^2\Phi}{d\omega dA |\omega \cdot \mathbf{n}_p|} = \frac{d^2\Phi}{d\omega dA \cos\theta_i} = \frac{d^2\Phi}{d\omega^\perp dA}, \quad (5.7)$$

where dA^\perp is the surface area around point \mathbf{p} projected on to the plane perpendicular to the direction ω . θ_i is the angle between the normal of the surface area at point \mathbf{p} and the direction ω . The relation between dA and dA^\perp is illustrated in Figure 5.6b.

Equation 5.7 show that the $\cos\theta_i$ factor can be absorbed by either the surface area dA , or the solid angle $d\omega$, in which case we are referring to the projected surface area dA^\perp , or the projected solid angle $d\omega^\perp$, respectively. We will use the definition with the projected solid angle in the derivations below.

An important observation is that the radiance is constant along an unobstructed ray in space. In addition, the sensor response in digital cameras is proportional to the radiance incident on it.

5.2.3 The BSDF

The Bidirectional Scattering Distribution Function (BSDF) is a mathematical description of the light-scattering properties of a surface. For two directions in the unit sphere at a surface point \mathbf{p} , it returns the ratio between the incoming and outgoing light ray's energy. It is defined as [74, 97, 79]:

$$f_s(\mathbf{p}, \omega_o, \omega_i) = \frac{dL_o(\mathbf{p}, \omega_o)}{dE(\mathbf{p}, \omega_i)} = \frac{dL_o(\mathbf{p}, \omega_o)}{L_i(\mathbf{p}, \omega_i) \cos\theta_i d\omega_i} \quad \left[\frac{1}{\text{sr}} \right]. \quad (5.8)$$

Chapter 5. BTDF Estimation with Radial Imaging

The BSDF is defined on the domain:

$$f_s : \mathcal{S}_i^2 \times \mathcal{S}_s^2 \rightarrow \mathbb{R},$$

where \mathcal{S}_i^2 and \mathcal{S}_s^2 are the incident and scattered spheres of directions, and $d\omega$ is the solid angle, as defined in equation 5.1.

The B(R/T)DF have the same definition as the BSDF:

$$f_{r/t}(p, \omega_o, \omega_i) = \frac{dL_o(p, \omega_o)}{dE(p, \omega_i)} = \frac{dL_o(p, \omega_o)}{L_i(p, \omega_i) \cos \theta_i d\omega_i} \left[\frac{1}{\text{sr}} \right], \quad (5.9)$$

but restricted domains. The domain of the BRDF is:

$$f_r : \mathcal{H}_i^2 \times \mathcal{H}_r^2 \rightarrow \mathbb{R},$$

where \mathcal{H}_i^2 and \mathcal{H}_r^2 are the incident and reflected hemispheres of directions. For BRDF $\mathcal{H}_r^2 = \mathcal{H}_i^2$ are the same hemispheres. They can either refer to the upward hemisphere \mathcal{H}_+^2 , or its complement \mathcal{H}_-^2 .

The BTDF is defined on the domain:

$$f_t : \mathcal{H}_i^2 \times \mathcal{H}_t^2 \rightarrow \mathbb{R}, \quad (5.10)$$

where \mathcal{H}_i^2 and \mathcal{H}_t^2 are the incident and transmitted hemispheres of directions, and where the transmitted hemisphere $\mathcal{H}_t^2 = -\mathcal{H}_i^2$ is the complement of the incident hemisphere.

Properties of the BRDF Physically based BRDF's have these basic properties:

1. Positivity $f_r(p, \omega_o, \omega_i) \geq 0$
2. Reciprocity $f_r(p, \omega_o, \omega_i) = f_r(p, \omega_i, \omega_o)$
3. Energy conservation $\int_{\mathcal{H}_t^2} f_r(p, \omega_o, \omega_i) \cos \theta_i d\omega_i \leq 1, \quad \forall \omega_o.$

The total energy of light reflected is less than or equal to the energy of the incident light.

These conditions are unique to reflection, since the BTDF does not always respect reciprocity.

Since the BRDF and the BTDF are both defined on the hemispheres of directions, we need two of each to compose the BSDF, one for each side of the surface. Therefore, purely reflective or transmissive surfaces are special cases of the BSDF.

5.3 The Light Transport Equation

We can rewrite equation 5.8 to express the differential outgoing radiance:

$$dL_o(p, \omega_o) = f_s(p, \omega_o, \omega_i) L_i(p, \omega_i) d\omega_i^\perp, \quad (5.11)$$

by also using equation 5.2 to absorb the $\cos\theta_i$ term into the projected solid angle $d\omega_i^\perp$.

By integrating equation 5.11 over all directions \mathcal{S}_i^2 of the unit sphere around point p , we get the scattering equation. This equation expresses the appearance of a surface, given an incident illumination description:

$$L_o(p, \omega_o) = \int_{\mathcal{S}^2} f_s(p, \omega_o, \omega_i) L_i(p, \omega_i) d\omega_i^\perp. \quad (5.12)$$

The reflectance equation refers to the corresponding equation for one-sided, opaque surfaces. It is derived by exchanging the BSDF with the BRDF, and integrating over the upward hemisphere of directions.

The Light Transport Equation (LTE), also known as the Rendering Equation, combines the scattered (equation 5.12) with the emitted radiance, and provides a complete specification of the light transport problem:

$$L_o(p, \omega_o) = L_e(p, \omega_o) + \int_{\mathcal{S}^2} f_s(p, \omega_o, \omega_i) L_i(p, \omega_i) d\omega_i^\perp. \quad (5.13)$$

It is customary to drop the subscripts for the radiance orientations because L_o and L_i can be expressed in terms of each other. We can therefore rewrite equation 5.13 as follows:

$$L_o(p, \omega_o) = L_e(p, \omega_o) + \int_{\mathcal{S}^2} f_s(p, \omega_o, \omega_i) L_o(\mathbf{x}_A(p, \omega_i), -\omega_i) d\omega_i^\perp. \quad (5.14)$$

where $\mathbf{x}_A(p, \omega_i)$ is a ray-casting function that selects the point on the closest surface in the




Figure 5.7 – Schematic of the rendering equation. The propagation operator \mathcal{G} transforms the far-field outgoing radiance into a local incident radiance field. The local scattering operator \mathcal{K} computes the outgoing radiance after a surface interaction, given the incoming radiance and the BSDF of the surface. Image taken from [9].

direction ω_i . We also have to reverse the direction of the second argument to L_o because the directions are expressed relative to the local point $\mathbf{x}_A(p, \omega_i)$. This form of the equation reveals the recursive nature of the problem.

5.3.1 Solution to the LTE

The solution to the LTE can elegantly be expressed using linear operators. A linear operator is a linear function whose domain is a vector space. In the case of light transport, it acts on radiance functions, and the result is another radiance function.

Light transport is considered to be an alternation of two steps. The first step is propagation, in which photons travel between scene surfaces in straight lines. The second step is scattering, and it describes the scattering of photons after a surface interaction. Figure 5.7 illustrates the two steps of the rendering equation.

Propagation operator

$$(\mathcal{G}L_o)(p, \omega_i) = \begin{cases} L_o(\mathbf{x}_A(p, \omega_i), -\omega_i) & \text{if } \mathbf{x}_A(p, \omega_i) \text{ is a valid surface point,} \\ 0 & \text{otherwise.} \end{cases} \quad (5.15)$$

The propagation operator \mathcal{G} represents the travel of light in straight lines through a fixed medium. It starts from the exitant radiance L_o that leaves other surfaces of the scene, and turns it into an incident radiance L_i , according to $L_i = \mathcal{G}L_o$. At the most basic level, it acts as a visibility mask.

Local scattering operator

$$(\mathcal{K}L_i)(\mathbf{p}, \omega_o) = \int_{S^2} f_s(\mathbf{p}, \omega_o, \omega_i) L_i(\mathbf{p}, \omega_i) d\omega_i^\perp. \quad (5.16)$$

The local scattering operator \mathcal{K} takes the incident radiance function L_i , and computes the exitant radiance $L_o = \mathcal{K}L_i$ that results from a single scattering operation. It operates on entire radiance functions, rather than on a single point \mathbf{p} , i.e., it operates over all surface points at once.

Light transport operator. By combining the effects of the propagation and the local scattering operators, we can assemble the light transport operator:

$$\mathcal{T} = \mathcal{K}\mathcal{G}. \quad (5.17)$$

This operator takes an exitant radiance function L_o and maps it into the exitant radiance function $\mathcal{T}L_o$. The application of \mathcal{T} results in a single scattering step.

Equilibrium radiance. We are interested in light transport computations when the scene radiance is at an equilibrium. This is known as the equilibrium radiance or the steady-state radiance. In this state, the exitant radiance is the sum of the emitted and the scattered radiance:

$$L_o = L_e + \mathcal{T}L_o. \quad (5.18)$$

The solution operator. We can find the light transport solution operator by inverting the equation for equilibrium radiance (equation 5.18):

$$(\mathcal{I} - \mathcal{T})L_o = L_e \quad \Rightarrow \quad L_o = (\mathcal{I} - \mathcal{T})^{-1}L_e, \quad (5.19)$$

where \mathcal{I} is the identity operator, and where the solution operator is:

$$\mathcal{S} = (\mathcal{I} - \mathcal{T})^{-1}. \quad (5.20)$$

We can now rewrite equation 5.18 by using \mathcal{S} :

$$L_o = \mathcal{S}L_e \quad (5.21)$$

The LTE solution operator has a physical interpretation. We can expand it into a Neumann series (a generalization of the geometric series):

$$\mathcal{S} = (\mathcal{I} - \mathcal{T})^{-1} = \sum_{i=0}^{\infty} \mathcal{T}^i = \mathcal{I} + \mathcal{T} + \mathcal{T}^2 + \dots \quad (5.22)$$

After this expansion, the equilibrium radiance can be rewritten as:

$$L_o = L_e + \mathcal{T}L_e + \mathcal{T}^2L_e + \dots \quad (5.23)$$

This expresses the outgoing radiance, L_o , as the sum of emitted radiance, and radiance that has been scattered once, twice, etc.

5.3.2 Surface form of the LTE

The solution operator from equation 5.20 provides us with a very elegant solution to the LTE. However, this hides the complexity of the BSDF models and the arbitrary scene geometries. Therefore, it is only possible to solve the LTE analytically in very constrained cases with simple BSDF's (e.g., Lambertian) and scene geometries. For general scenes we have to use numerical integration. That is the reason why researchers have invested heavily in ray tracing algorithms and Monte Carlo integration.

Equation 5.14 can be solved with Monte Carlo integration by sampling from a distribution on the sphere of directions and casting rays. A more elegant way to evaluate the integrand is to sample points on scene surfaces, evaluate the coupling between the points, and trace rays to evaluate their visibility. To rewrite equation 5.14 from an integral over directions to an integral over area, we need to define some substitutions. We will use Figure 5.8 to aid our discussion. There we can see that p' is a point on a reflecting surface, and we would like to compute the radiance leaving in the direction of point p , and we also have an incident radiance from a surface that point p'' lies on. We start by redefining the directions:

$$\omega_i = \widehat{p' - p''} = p'' \rightarrow p', \quad \omega_o = \widehat{p - p'} = p' \rightarrow p. \quad (5.24)$$




Figure 5.8 – The three-point form of the LTE. Incident radiance from the surface around p'' gets scattered by the surface around p' , which in turn converts it to outgoing radiance towards p . This example scene is used to illustrate the conversion of the LTE from an integral over directions on the unit sphere, to an integral over points on scene surfaces. This is a modified version of a figure from [79].

If points p' and p are mutually visible, we can now define the exitant radiance from p' to p as:

$$L(p' \rightarrow p) = L(p', \omega_o), \quad (5.25)$$

and the BSDF as:

$$f_s(p'' \rightarrow p' \rightarrow p) = f_s(p', \omega_o, \omega_i). \quad (5.26)$$

Furthermore, we introduce a geometric coupling term, also known as geometric kernel, that includes a visibility term between surface points:

$$G(p \leftrightarrow p') = V(p \leftrightarrow p') \frac{|\cos \theta| |\cos \theta'|}{\|p - p'\|^2} \quad (5.27)$$

We can finally substitute equations 5.25, 5.26, and 5.27 into equation 5.14, and convert the integral to operate over surface area instead of solid angle by using 5.2, to get the surface form of the LTE:

$$L(p' \rightarrow p) = L_e(p' \rightarrow p) + \int_A f_s(p'' \rightarrow p' \rightarrow p) L(p'' \rightarrow p') G(p'' \leftrightarrow p') dA(p''), \quad (5.28)$$

where A includes the whole surface area of the scene, i.e., the surface area of all scene objects.

5.3.3 Discretized surface form of the LTE

If we can adequately discretize the surface areas of the scene objects, we can rewrite the surface form of the LTE as a sum instead of an integral:

$$L(\mathbf{p}' \rightarrow \mathbf{p}) = L_e(\mathbf{p}' \rightarrow \mathbf{p}) + \sum_A f_s(\mathbf{p}'' \rightarrow \mathbf{p}' \rightarrow \mathbf{p}) L(\mathbf{p}'' \rightarrow \mathbf{p}') G(\mathbf{p}'' \leftrightarrow \mathbf{p}') A(\mathbf{p}''). \quad (5.29)$$

This formulation will produce correct results only in the case where the facets are small enough such that the BSDF and the outgoing radiances do not change significantly inside the surface facets. Furthermore, the discrete approximation of the geometric coupling is correct only in the case where the surface facets are infinitesimal and distant from each other [71, 19].

5.3.4 Analytical solution to the LTE for a simple scene

Let us now analyze the scene from Figure 5.2. We discretize the cylinder in N facets, which can be seen in Figure 5.9. In this figure, the facets are as square a shape as possible and have equal areas. Having equal areas is not necessary, since the surface area is taken into account in the equations, but it does help with having a more elegant and memory efficient implementation.

Radiance from facet \mathbf{p}_j to facet \mathbf{p}_i . We will start with the discretized surface form of the LTE (equation 5.29), since our cylinder is discretized into facets:

$$L(\mathbf{p}_j \rightarrow \mathbf{p}_i) = L_e(\mathbf{p}_j \rightarrow \mathbf{p}_i) + \sum_{k=1}^N f_r(\mathbf{p}_k \rightarrow \mathbf{p}_j \rightarrow \mathbf{p}_i) L(\mathbf{p}_k \rightarrow \mathbf{p}_j) G(\mathbf{p}_k \leftrightarrow \mathbf{p}_j) A(\mathbf{p}_k), \quad (5.30)$$

where N is the total number of cylinder facets. All facets are visible to each other, except the ones that are collinear. However, even in this case there is no need for the visibility operator, since the polar angles θ for those facets are $\pi/2$, so the cosines in G yield zero. Furthermore, we use the BRDF f_r instead of the BSDF f_s , since we consider the cylinder to be opaque.

The cylinder is not radiant, which makes the term $L_e(\mathbf{p}_j \rightarrow \mathbf{p}_i) = 0$. However, we will assume that the facets do radiate light, thus $L_e(\mathbf{p}_j \rightarrow \mathbf{p}_i) > 0$. We can do this by assuming that L_e is the first bounce radiance, i.e., the radiance that leaves the sample and bounces off of the cylinder facets only once. The radiance $L(\mathbf{p}_k \rightarrow \mathbf{p}_j)$ is the equilibrium outgoing radiance, after an infinite number of light bounces inside the cylinder.




Figure 5.9 – The cylinder from Figure 5.2, discretized into 5184 facets that are as square as possible and have equal areas.

5.3.5 Radiance of a single facet

We can remove the sum from equation 5.30 by rewriting it in matrix form:

$$L(p_j \rightarrow p_i) = L_e(p_j \rightarrow p_i) + \mathbf{f}_r^{p_j \rightarrow p_i^T} \mathbf{G}^{\rightarrow p_j} \mathbf{I}^{\rightarrow p_j}, \quad (5.31)$$

where:

$$\begin{aligned} \mathbf{f}_r^{p_j \rightarrow p_i} &= (f_r(p_1 \rightarrow p_j \rightarrow p_i), f_r(p_2 \rightarrow p_j \rightarrow p_i), \dots, f_r(p_N \rightarrow p_j \rightarrow p_i))^T, \\ \mathbf{G}^{\rightarrow p_j} &= \text{diag}(G(p_1 \leftrightarrow p_j) A(p_1), G(p_2 \leftrightarrow p_j) A(p_2), \dots, G(p_N \leftrightarrow p_j) A(p_N)), \\ \mathbf{I}^{\rightarrow p_j} &= (L(p_1 \rightarrow p_j), L(p_2 \rightarrow p_j), \dots, L(p_N \rightarrow p_j))^T. \end{aligned}$$

We can express equation 5.31 more compactly by using a vector instead of a diagonal matrix for $\mathbf{G}^{\rightarrow p_j}$, and taking the Hadamard (elementwise) product of the vectors:

$$L(p_j \rightarrow p_i) = L_e(p_j \rightarrow p_i) + \mathbf{f}_r^{p_j \rightarrow p_i^T} (\mathbf{g}^{\rightarrow p_j} \odot \mathbf{I}^{\rightarrow p_j}). \quad (5.32)$$

Radiance from facet p_j to all facets. We can easily extend the previous equation 5.32 to account for the radiance reflected by a single facet to all other facets. The last term in the

parenthesis is the irradiance that facet p_j receives from all other cylinder facets, thus it remains unchanged. However, we now have to extend the L_e term into a vector and the BRDF term into a matrix to account for all outgoing facets:

$$\mathbf{l}^{p_j \rightarrow} = \mathbf{l}_e^{p_j \rightarrow} + \mathbf{F}_r^{p_j} \mathbf{G}^{\rightarrow p_j} \mathbf{l}^{\rightarrow p_j} \quad (5.33)$$

$$= \mathbf{l}_e^{p_j \rightarrow} + \mathbf{F}_r^{p_j} (\mathbf{g}^{\rightarrow p_j} \odot \mathbf{l}^{\rightarrow p_j}), \quad (5.34)$$

where:

$$\mathbf{l}_e^{p_j \rightarrow} = (L_e(p_j \rightarrow p_1), L_e(p_j \rightarrow p_2), \dots, L_e(p_j \rightarrow p_N))^T, \\ \mathbf{F}_r^{p_j} = (\mathbf{f}_r^{p_j \rightarrow p_1}, \mathbf{f}_r^{p_j \rightarrow p_2}, \dots, \mathbf{f}_r^{p_j \rightarrow p_N})^T.$$

5.3.6 Radiance between all facets

The final step is to compute the radiance from all facets to all facets. Expanding on the previous equations, we can write the following relation:

$$\mathbf{l} = \mathbf{l}_e + \mathbf{F}_r \mathbf{G} \mathbf{l} \\ = \mathbf{l}_e + \mathbf{F}_r (\mathbf{g} \odot \mathbf{l}), \quad (5.35)$$

where $\mathbf{F}_r \in \mathbb{R}^{N^2 \times N^2}$ has a complex structure, which can be seen in Appendix 5.A, and:

$$\mathbf{l} = \begin{bmatrix} L(p_1 \rightarrow p_1) \\ L(p_1 \rightarrow p_2) \\ \vdots \\ L(p_1 \rightarrow p_N) \\ L(p_2 \rightarrow p_1) \\ \vdots \\ L(p_N \rightarrow p_N) \end{bmatrix}, \quad \mathbf{l}_e = \begin{bmatrix} L_e(p_1 \rightarrow p_1) \\ L_e(p_1 \rightarrow p_2) \\ \vdots \\ L_e(p_1 \rightarrow p_N) \\ L_e(p_2 \rightarrow p_1) \\ \vdots \\ L_e(p_N \rightarrow p_N) \end{bmatrix}, \quad \mathbf{g} = \begin{bmatrix} G(p_1 \leftrightarrow p_1) A(p_1) \\ G(p_1 \leftrightarrow p_2) A(p_1) \\ \vdots \\ G(p_1 \leftrightarrow p_N) A(p_1) \\ G(p_2 \leftrightarrow p_1) A(p_2) \\ \vdots \\ G(p_N \leftrightarrow p_N) A(p_N) \end{bmatrix}, \quad \mathbf{G} = \begin{bmatrix} \mathbf{g}^T \\ \mathbf{g}^T \\ \vdots \\ \mathbf{g}^T \\ \mathbf{g}^T \\ \vdots \\ \mathbf{g}^T \end{bmatrix},$$

$$\mathbf{l} \in \mathbb{R}^{N^2}, \quad \mathbf{l}_e \in \mathbb{R}^{N^2}, \quad \mathbf{g} \in \mathbb{R}^{N^2}, \quad \mathbf{G} \in \mathbb{R}^{N^2 \times N^2}.$$

Equation 5.35 is the discrete surface form of the equilibrium radiance for our simple scene. It is very similar to the equilibrium radiance equation in operator form, equation 5.18 ($L = L_e + \mathcal{T}L$).

In the forward rendering problem, we are given L_e together with the scene description, and we want to compute L . The solution to equation 5.35 is:

$$L = \mathcal{S}L_e = (\mathcal{I} - \mathcal{T})^{-1} L_e \quad (5.36)$$

$$\mathbf{l} = (\mathbf{I} - \mathbf{F}_r \mathbf{G})^{-1} \mathbf{l}_e. \quad (5.37)$$

In our case, we are trying to solve the inverse problem, where we observe the equilibrium radiance L , and we want to compute L_e :

$$\begin{aligned} L_e &= \mathcal{S}^{-1}L = (\mathcal{I} - \mathcal{T})L \\ \mathbf{l}_e &= (\mathbf{I} - \mathbf{F}_r \mathbf{G})\mathbf{l}. \end{aligned} \quad (5.38)$$

Note that we consider L_e to be the radiance that the facets scatter after the first bounce, and not the radiance emitted by the sample (or, traditionally, by the light source), i.e., we move the computation one bounce forward.

5.3.7 Lambertian cylinder

If the cylinder is coated with a Lambertian material, we can directly use equation 5.38. What is interesting is that we can further simplify it due to the very simple BRDF and radiance expressions.

Lambertian BRDF and radiance

The Lambertian (diffuse) BRDF is defined as:

$$f_r(p_k \rightarrow p_j \rightarrow p_i) = \frac{\rho}{\pi}. \quad (5.39)$$

Since we are dealing with a cylinder whose surface is homogenous, we can assume that every facet of the cylinder has the same reflectance ρ . This is not a difficult constraint to overcome, and we will see how we can handle it in equation 5.43.

Each radiance measurement can be considered as the response of a hypothetical sensor placed somewhere in the scene. Therefore, we will assume that the surface of the sensor is also part

Chapter 5. BTDF Estimation with Radial Imaging

of the scene geometry, and we will denote the lens aperture by p_0 .

The diffuse radiance leaving a facet has the same magnitude in every direction. That same radiance is also measured by the image sensor:

$$L(p_k \rightarrow p_j) = L(p_k \rightarrow p_0). \quad (5.40)$$

By plugging equations 5.39 and 5.40 in equation 5.30, the radiance of the Lambertian cylinder can be expressed as:

$$L(p_j \rightarrow p_i) = L_e(p_j \rightarrow p_i) + \frac{\rho}{\pi} \sum_{k=1}^N L(p_k \rightarrow p_0) G(p_k \leftrightarrow p_j) A(p_k). \quad (5.41)$$

Solution to radiance from all facets to the image sensor - Lambertian cylinder

With a Lambertian cylinder we can modify equation 5.35 to express the radiance from all facets to facet p_0 (the image sensor) as follows:

$$\mathbf{I}^{\rightarrow p_0} = \mathbf{I}_e^{\rightarrow p_0} + \frac{\rho}{\pi} \mathbf{G}^L \mathbf{I}^{\rightarrow p_0}, \quad (5.42)$$

$$\begin{bmatrix} L(p_1 \rightarrow p_0) \\ L(p_2 \rightarrow p_0) \\ \vdots \\ L(p_N \rightarrow p_0) \end{bmatrix} = \begin{bmatrix} L_e(p_1 \rightarrow p_0) \\ L_e(p_2 \rightarrow p_0) \\ \vdots \\ L_e(p_N \rightarrow p_0) \end{bmatrix} + \frac{\rho}{\pi} \begin{bmatrix} \mathbf{g}^{\rightarrow p_1^T} \\ \mathbf{g}^{\rightarrow p_2^T} \\ \vdots \\ \mathbf{g}^{\rightarrow p_N^T} \end{bmatrix} \begin{bmatrix} L(p_1 \rightarrow p_0) \\ L(p_2 \rightarrow p_0) \\ \vdots \\ L(p_N \rightarrow p_0) \end{bmatrix},$$

$$\mathbf{I}^{\rightarrow p_0} \in \mathbb{R}^N, \quad \mathbf{I}_e^{\rightarrow p_0} \in \mathbb{R}^N, \quad \mathbf{G}^L \in \mathbb{R}^{N \times N}.$$

We now use a compact projected solid angle matrix \mathbf{G}^L and reuse the same $\mathbf{I}^{\rightarrow p_0}$ for all facets. The simplified structure of the solution operator for the inverse rendering problem, in the case of a Lambertian cylinder, can be expressed as follows:

$$\mathbf{I}_e^{\rightarrow p_0} = \left(\mathbf{I} - \frac{\rho}{\pi} \mathbf{G}^L \right) \mathbf{I}^{\rightarrow p_0}. \quad (5.43)$$

The steady-state radiance, $\mathbf{I}^{\rightarrow p_0}$, is measured by the image sensor, the projected solid angle, \mathbf{G}^L ,




Figure 5.10 – The outgoing BRDF lobes of Spectralon for incident polar angles θ_i of (a) $\theta_i = 0^\circ$ and (b) $\theta_i = 75^\circ$ (represented by the vertical pink lines). The outgoing lobe for $\theta_i = 0^\circ$ is close to hemispherical (Lambertian BRDF). However, at steep incident angles ($\theta_i = 75^\circ$) it has an obvious specular component. Images taken from the BRDF visualization tool Tekari.³

is governed by the geometry of the scene, and ρ is the reflectance of the Lambertian coating. If the reflectance varies across the surface of the cylinder, it can be expressed as a diagonal matrix filled with the reflectance of each facet. We do, however, assume that the reflectance inside the facets is constant.

5.3.8 Non-Lambertian cylinder

The Lambertian case discussed above is an interesting theoretical experiment. However, no known material has a completely diffuse BRDF. In a more realistic case, where the coating of the cylinder is non-Lambertian, we cannot directly use equation 5.38. The problem is that we cannot measure the exitant radiance at equilibrium of each facet, towards all other facets (vector \mathbf{l}). Even if we could measure \mathbf{l} , the solution matrix would be too large for any practical application with a decent number of facets ($N^2 \times N^2$).

If we have a closer look at the solution operator from equation 5.20, we can see that in the discrete surface form of the LTE it is a light transport matrix (LTM). The LTM encapsulates the light paths inside our scene, i.e., the influence that the light source radiance has on the steady-state exitant radiance captured by the camera. The light transport inside the scene can be completely described by a matrix because it is a linear operation [76, 88].

In this work we are interested in solving the inverse rendering problem. We know the captured image L and the scene geometry, which is represented by the geometric coupling G (equation 5.27), and we would like to retrieve the radiance distribution L_e of the sample, which we can also consider as the light source.

³<https://rgl.epfl.ch/software/Tekari>

Having that in mind, we would like to generate a light transport matrix that provides a bijective mapping between the radiance of the sample issued in the direction of every cylinder facet, and the equilibrium radiance of every facet captured by the camera. This will result in a matrix of size $N \times N$.

5.4 Light transport matrix

We can imagine that the light leaving the sample inside the cylinder is collimated in a certain direction ω . This corresponds to a radiance $L(p_s, \omega)$, where p_s represents the center of the sample. In an ideal world, this unidirectional radiance irradiates exactly one facet of the cylinder. The facet then reflects this light according to its BRDF, and after an infinite number of interreflections, the cylinder reaches a steady state. In this state, there is a bijective relationship between $L(p_s, \omega)$, and the irradiance distribution E_{cyl} on the cylinder.

The steady-state radiance flowing from the facets to the pixels of the image sensor creates an irradiance distribution $E_{cam}(x, y; \omega)$ on the sensor, where x and y are the coordinates of the image pixels. This irradiance distribution is unique for the considered incidence direction. If we change the incident radiance, for example we multiply it by a factor of two, we do not modify the distribution, but just its magnitude (multiplied by a factor of two), since it is linear.

Now we can assume that the incident radiance from the sample to the facets has a certain angular distribution $L(p_s, \omega)$ that is not collimated. This radiance is the sum of several collimated radiances, each one producing its irradiance distribution $E_{cam}(x, y; \omega)$ on the image sensor. Since they are independent from each other (there is no interference), they simply add up. We can represent this more formally:

$$\begin{bmatrix} L(p_1 \rightarrow p_0) \\ L(p_2 \rightarrow p_0) \\ \vdots \\ L(p_N \rightarrow p_0) \end{bmatrix} = \begin{bmatrix} t_{1,1} & t_{1,2} & \cdots & t_{1,N} \\ t_{2,1} & & \ddots & \vdots \\ \vdots & & & \\ t_{N,1} & & & t_{N,N} \end{bmatrix} \begin{bmatrix} L(p_s \rightarrow p_1) \\ L(p_s \rightarrow p_2) \\ \vdots \\ L(p_s \rightarrow p_N) \end{bmatrix}. \quad (5.44)$$

We can rewrite this equation in terms of the initial irradiance of each facet by the sample E_0 , since there exists a bijective relationship between $E_0(p_i)$ and $L(p_s \rightarrow p_i)$:

$$E_0(p_i) = L(p_s \rightarrow p_i) G(p_s \leftrightarrow p_i) A(p_s). \quad (5.45)$$

We can now express equation 5.44 in terms of the initial facet irradiance distribution E_0 :

$$\begin{bmatrix} L(p_1 \rightarrow p_0) \\ L(p_2 \rightarrow p_0) \\ \vdots \\ L(p_N \rightarrow p_0) \end{bmatrix} = \begin{bmatrix} t_{1,1} & t_{1,2} & \cdots & t_{1,N} \\ t_{2,1} & \ddots & & \vdots \\ \vdots & & & \\ t_{N,1} & & & t_{N,N} \end{bmatrix} \begin{bmatrix} E_0(p_1) \\ E_0(p_2) \\ \vdots \\ E_0(p_N) \end{bmatrix}. \quad (5.46)$$

If only one facet is irradiated by the sample, there will be a single $E_0(p_k)$ that is non-zero. The radiance distribution that this $E_0(p_k)$ produces is given by the k^{th} column of the matrix:

$$\begin{bmatrix} L(p_1 \rightarrow p_0) \\ L(p_2 \rightarrow p_0) \\ \vdots \\ L(p_N \rightarrow p_0) \end{bmatrix} = \begin{bmatrix} \cdots & t_{1,k} & \cdots & \cdots \\ \cdots & t_{2,k} & \cdots & \cdots \\ \vdots & \vdots & \vdots & \vdots \\ \cdots & t_{N,k} & \cdots & \cdots \end{bmatrix} \begin{bmatrix} 0 \\ E_0(p_k) \\ \vdots \\ 0 \end{bmatrix}. \quad (5.47)$$

We can see that there is a bijective, linear relationship between E_0 and L . In equation 5.45 we also saw that there is a linear relationship between the incoming radiance function $L(p_s, \omega)$ and E_0 . Therefore, there is a linear relationship between the radiance distribution emitted by the sample $L(p_s, \omega)$ and the radiance measured by the camera L .

5.4.1 Computing the light transport matrix

The light transport matrix describes how does an emitted radiance affects the final image. In our case, we are interested in how the irradiance of each facet influences the equilibrium radiance captured by the camera. It is the light reflected towards the camera after an infinite number of light bounces, after the light took all possible light paths. The solution operator for the LTE also provides us with the same concept.

We saw that we can apply the solution operator to solve for infinite bounces in the Lambertian case of our scene. We can slightly modify equation 5.42 to express it in terms of the initial facet irradiance vector \mathbf{e}_0 , instead of the “emitted” radiance by the facets:

$$\begin{aligned} \mathbf{I}^{\rightarrow p_0} &= \frac{\rho}{\pi} \left(\mathbf{I} - \frac{\rho}{\pi} \mathbf{G}^L \right)^{-1} \mathbf{e}_0 \\ &= \frac{\rho}{\pi} \left(\mathbf{I} + \frac{\rho}{\pi} \mathbf{G}^L + \left(\frac{\rho}{\pi} \mathbf{G}^L \right)^2 + \cdots \right) \mathbf{e}_0 \end{aligned} \quad (5.48)$$

In the non-Lambertian case, the cylinder BRDF complicates the matters, and thus we cannot easily compute the solution operator since we do not have access to the radiance scattered

by each facet towards all other facets. Therefore, we have to resort to computing each light bounce explicitly:

$$\begin{aligned} \mathbf{I}^{\rightarrow p_0} &= \mathbf{T} \mathbf{e}_0 \\ &= (\mathbf{T}_1 + \mathbf{T}_2 + \mathbf{T}_3 + \dots) \mathbf{e}_0 \end{aligned} \quad (5.49)$$

The light transport matrix \mathbf{T} can be computed with ray tracing. We can assume that the sample irradiates only facet k . By tracing the rays for a single bounce, we can compute the corresponding column k of the \mathbf{T}_m matrix, where m is the number of bounces. We then execute another bounce, and populate column k of matrix \mathbf{T}_{m+1} . After computing column k for a predetermined number of bounce, we repeat the steps for facet $k + 1$, until all facets have been irradiated individually.

Computing the light transport matrix is expensive. That is why parallelizing the computation is interesting. In their dual photography paper, Sen et al. [89] were able to reduce the number of captured images by isolating groups of projector pixels that do not affect the same camera pixels when projected together. This is possible when the scene has isolated pockets that do not interact with other scene surfaces. Unfortunately, in our simple scene, all facets are visible from, and influence all other facets.

Fortunately, due to the simplicity of the scene we can still reduce the computational load. The scene is completely symmetric, since the cylinder is rotationally symmetric and the sample and camera are positioned along the center axis of the cylinder. Therefore, it is enough to fill the light transport matrix for the facets along a single phi angle and all heights. With simple circular shifts of this data we can fill the light transport matrix for all other facets, due to the rotational symmetry of the cylinder.

In the first light bounce, the LTM scatters the initial irradiance of the facets towards the camera as radiance. Scattering the irradiance in the direction of the camera, can be expressed by multiplying the irradiance with the BRDF of the cylinder. The incident direction of the BRDF is from the sample p_s towards a facet p_i , and an outgoing direction is from the facet p_i towards the camera p_0 . This also means that there are no interactions between the facets. We can express this as a diagonal matrix that contains only the BRDF of the cylinder coating as its diagonal elements:

$$\mathbf{T}_1 = \text{diag}(f_r(p_s \rightarrow p_1 \rightarrow p_0), f_r(p_s \rightarrow p_2 \rightarrow p_0), \dots, f_r(p_s \rightarrow p_N \rightarrow p_0)). \quad (5.50)$$

For the second bounce, the radiance leaving facet p_i towards the camera is the sum of the radiance sent by each facet towards facet p_i , which is then scattered towards the camera.

Analogous to scattering the initial facet irradiance towards the camera in the single bounce, we multiply the irradiance by $f_r(p_s \rightarrow p_j \rightarrow p_i)$. We then calculate the irradiance of facet p_i by all other facets by taking into account the geometric coupling and the surface areas of those facets. Finally, facet p_i reflects this irradiance towards the camera. In this step we are only interested in the second bounce radiance, thus we will not have the first bounce elements on the LTM diagonal. Since now the facets interact with each other, we have a dense matrix:

$$\mathbf{T}_2(i, j) = f_r(p_j \rightarrow p_i \rightarrow p_0) G(p_j \rightarrow p_i) f_r(p_s \rightarrow p_j \rightarrow p_i).$$

Note that in the previous equation we use a unidirectional $G(p_j \rightarrow p_i)$, instead of the bidirectional $G(p_j \leftrightarrow p_i)$ that we used up until now. We define it as follows:

$$G(p_j \rightarrow p_i) = G(p_j \leftrightarrow p_i) A(p_j)$$

For the third and further bounces, we repeat the second bounce one more time. We now have the light reflecting from all facets towards facet p_j . This is the irradiance of facet p_j from all other facets, which then get reflected towards facet p_i . This is represented by the sum over all facets N . The radiance from each facet irradiates p_i , which then reflects this irradiance towards the camera p_0 :

$$\mathbf{T}_3(i, j) = f_r(p_j \rightarrow p_i \rightarrow p_0) G(p_j \rightarrow p_i) \sum_{k=1}^N f_r(p_k \rightarrow p_j \rightarrow p_i) G(p_k \rightarrow p_j) f_r(p_s \rightarrow p_k \rightarrow p_j).$$

Further bounces can then be recursively added until a desired number has been reached. Figure 5.11 shows the computational schematic for a column of the LTM. The discretization is symbolic and it assumes only three facets. Wherever an x appears in the subscripts, it should be replaced with the corresponding column number of the matrix. The blocks of different color represent the different parts of the computation. The matrices in the bottom row show the intermediate results of the computation. All operations between the matrices are Hadamard products.

The computation starts from the left side, and progresses towards the right. It starts with the radiance that the sample issues towards the facets $\mathbf{I}^{p_s \rightarrow}$. As explained above, each cycle of computation produces a single column of the LTM matrix. Therefore, $\mathbf{I}^{p_s \rightarrow}$ has a single one and $N - 1$ zero entries, such that a single facet is illuminated by the sample. This radiance, when it is multiplied with the geometric coupling and the area of the sample $\mathbf{g}^{p_s \rightarrow}$, is received by the facets as the initial irradiance \mathbf{e}_0 . It then gets scattered towards all other facets when multiplied




Figure 5.11 – An iterative computation of a single column of the light transport matrix. The first part, with green background, turns the incoming radiance from the sample towards each facet into an outgoing radiance of each facet towards all other facets. The middle part, with red background, accounts for the light bounces inside the cylinder. It turns the incoming radiance from all facets towards one facet, and then distributes it again towards all other facets, according to the BRDF of the cylinder. This is done for all cylinder facets and it can be repeated as many times as necessary to account for multiple light bounces. The final step, with blue background, is responsible for turning the incoming radiance from all facets towards a single facet, and then relaying it to the camera, i.e., to the corresponding pixel of the virtual image. This is also done for all facets. All operations are additions and multiplications, which can be computed efficiently on the GPU.

with the BRDF, that has the sample as an incoming direction, and the corresponding outgoing facet directions. In the red block the facets receive the radiance reflected by all other facets as irradiance when multiplied with the geometric coupling and the area of the corresponding reflecting facet. This irradiance is then reflected back towards all facets when applying the BRDF. The BRDF is represented as a 3D tensor, since it has the reflecting facet as the first dimension, the incoming facet as the second dimension, and the outgoing facet as the third dimension. To reduce the result to a matrix, we sum over the second dimension (incoming facet) and permute the dimensions such that the outgoing facet is now the first dimension, and the reflecting facet is the second. This block can be recursively executed as many times as necessary. A single execution will result in three light bounces, two executions will result in four light bounces, etc. Finally, the blue block takes the incoming radiance from all facets as irradiance, and scatters it towards the camera. At the end, we reduce the second dimension (incoming facet), and we are left with the radiance scattered by the reflecting facets towards the camera as a column vector. This process is then repeated by selecting a different facet to be illuminated by the sample.

5.5 Algorithm

Our novel approach for estimating a sample's BTDF is based on the schematic from Figure 5.2. An overview of our algorithm can be found in Figure 5.12. It starts with a calibration to get the sample, cylinder and camera properly aligned. We also have to capture the intensity of the light source and the lens vignetting. The calibration is an offline procedure.

Once we calibrate the system, we can then insert the sample and capture a set of images with different exposure times of the cylinder inside. We preprocess these images and assemble them in an High-Dynamic Range (HDR) image. For the simulations, we do not need the pre-processing step, because the image is already in HDR. There is also no need for the calibration step, since the sample, cylinder and camera are perfectly aligned, the intensity of the light source is known, and there is no vignetting.

We then create a virtual cylinder with the same height and radius as the real cylinder, and discretize it into facets. Since we have the geometric calibration of the camera, we can project the discretized cylinder onto the HDR image. The projected facets are used to classify the image pixels into different virtual pixels, and the virtual pixels give us the radiance of the facets. This allows us to create the *virtual image* from the HDR image.

Now we are ready to compute the first bounce irradiance of each facet E_0 . To cancel the infinite interreflections inside the cylinder, we invert the light transport matrix, and apply it to the equilibrium radiance of the facets. In the final step, we compute the BTDF of the sample from E_0 .

5.5.1 Calibration

The calibration procedure is divided in geometric and radiometric. The geometric calibration is used to align the sample and cylinder to the camera, by matching the projection of the real cylinder onto the camera image to the projection of an ideal cylinder. The center of the sample should be aligned with the center of the cylinder, and with the principal point of the image. This alignment is important, because the LTM is computed for the sample, cylinder and camera at fixed relative positions. Deviations from those positions will introduce errors in the computed BTDF. We perform the calibration manually: we adjust the camera and cylinder positions, take a photo, project the expected positions of the cylinder and sample on the captured image, and we repeat the process until we align them properly. This procedure is executed only once, and is valid as long as the scene geometry does not change.

For the radiometric calibration we capture an image of the unobstructed light beam of the light source. Its role is to provide the scale for the computed BTDF, thus turning relative to absolute values. The radiometric calibration is also valid as long as the light source intensity remains constant.




Figure 5.12 – Flowchart of our algorithm for BTDF estimation. The calibration procedure is divided in geometric and radiometric. These are done offline, and are valid as long as the scene geometry and light source intensity remain constant. In the preprocessing step we merge a set of images taken with different exposures into a single HDR image. The HDR image can also be the output of a rendering software. The cylinder facets of a virtual cylinder are then projected onto the HDR image to produce a virtual image, which represents the equilibrium radiance of the facets. After reversing the infinite light bounces of the virtual image, we can calculate the BTDF of the measured sample. The preprocessing, virtual image computation, and the calculation of the BTDF are done online.



Figure 5.13 – In this figure we can see (a) a captured image, (b) a modeled background image, and (c) the captured image after normalizing its pixel intensities with the modeled background image. The modeled background image is produced after capturing a uniform white screen and fitting a 4th order polynomial to it to remove noise and other image artifacts.

5.5.2 Preprocessing

Vignetting correction. Vignetting is the radial falloff of intensity towards the edges and corners of an image. It can be caused by various sources [31]:

- Natural vignetting: due to geometric optics, also known as the cosine⁴ law.
- Pixel vignetting: light striking a photon well can be partially occluded by its sides.
- Optical vignetting: light paths blocked inside the lens body by the lens diaphragm. This is a function of aperture size, and can be reduced by decreasing the aperture.
- Mechanical vignetting: light paths blocked by lens filters, hoods, etc.

To remove the vignetting of our camera and lens setup, we captured an image of a uniform white screen. We then fit a 4th order polynomial model to get rid of noise, scratches, smudges, and other image artifacts. This creates a correction image that we can apply to normalize the captured images by dividing their pixel intensities with this image's.

An example captured image can be seen in Figure 5.13a. The modeled background image is visible in Figure 5.13b. The final image after normalizing with the modeled background image can be seen in Figure 5.13c.

HDR image. The steady-state radiance values inside the cylinder exhibit a high dynamic range. It is higher than the dynamic range of the camera that we have used. Therefore, we needed to capture several images with different exposure times and assemble them in a single HDR image. We captured linear images, i.e., the sensor intensities had linear response to incident radiance, and we applied vignetting correction to the captured images. Also, our images are geometrically registered, because our setup is static and there are no moving objects inside the scene.

Chapter 5. BTDF Estimation with Radial Imaging

To obtain the HDR image, we use the algorithm by Debevec and Malik [18]. In a linear imaging system like ours, and assuming that the pixel irradiance is constant due to the steady-state radiance assumption of the scene:

$$Z_{ij} = E_i \Delta t_j, \quad (5.51)$$

where Z_{ij} is the pixel value of pixel i in image j , E_i is the pixel irradiance, and Δt_j is the exposure duration of image j . From the previous equation we can recover the pixel irradiance with the following equation:

$$\ln E_i = \ln Z_{ij} - \ln \Delta t_j. \quad (5.52)$$

The algorithm uses the pixel values of each captured image for a given pixel location i , and gives higher weight to exposures in which the pixel's value is closer to the middle of the pixel value range. This increases robustness, as it removes noise and other artifacts that might appear towards the extremes of the pixel value range, such as non-linear sensor response and saturated pixel values:

$$\ln E_i = \frac{\sum_{j=1}^P w(Z_{ij}) (\ln Z_{ij} - \ln \Delta t_j)}{\sum_{j=1}^P w(Z_{ij})}, \quad (5.53)$$

with the weighting scheme:

$$w(z) = \begin{cases} z - Z_{min} & \text{for } z \leq \frac{1}{2}(Z_{min} + Z_{max}), \\ Z_{max} - z & \text{for } z > \frac{1}{2}(Z_{min} + Z_{max}). \end{cases} \quad (5.54)$$

Undoing the radial and tangential distortion of the image. As a final step in the image preprocessing we undo lens distortions of the HDR image. We calibrated the camera using the method that we presented in Chapter 2. We performed only a geometric calibration of the camera.

It is important to correct these distortions because in the later steps we would like to establish correspondences between the captured image and the actual 3D points on the cylinder. These distortions introduce nonlinear deviations to the projections of the cylinder 3D points from their expected locations on the image.

For calibrating the camera we use at least 20 images of the calibration chart. We then recover 3 radial and 2 tangential distortion coefficients. The radial distortion is radially symmetric, due to the symmetry of the camera lens. The tangential distortion happens because of imperfect centering of the lens components and other manufacturing defects in a compound lens. For more details, please refer to Brown [14]. The following equation is used to rectify the image pixel locations:

$$\begin{bmatrix} x_u \\ y_u \end{bmatrix} = \begin{bmatrix} x_d + \bar{x}(k_1 r^2 + k_2 r^4 + k_3 r^6 + \dots) + [p_1(r^2 + 2\bar{x}^2) + 2p_2\bar{x}\bar{y}](1 + p_3 r^2 + p_4 r^4 + \dots) \\ y_d + \bar{y}(k_1 r^2 + k_2 r^4 + k_3 r^6 + \dots) + [2p_1\bar{x}\bar{y} + p_2(r^2 + 2\bar{y}^2)](1 + p_3 r^2 + p_4 r^4 + \dots) \end{bmatrix}, \quad (5.55)$$

where k and p are the radial and tangential distortion coefficients, x_d and y_d are the pixel coordinates of the distorted image, x_u and y_u are the pixel coordinates of the undistorted image, x_0 and y_0 are the coordinates of the distortion center, $\bar{x} = x_d - x_0$, $\bar{y} = y_d - y_0$, and $r = \sqrt{\bar{x}^2 + \bar{y}^2}$. The distortion is then undone by interpolating the image to the updated pixel grid.

5.5.3 Virtual image

To generate the virtual image, we use the preprocessed HDR image. Alternatively, when we use a rendered simulation image, we do not need the preprocessing steps, because the image is already in HDR, there is no vignetting or distortion artifacts, and the sample, cylinder and camera are perfectly aligned.

We create a virtual image from the camera image mainly to decrease the number of pixels, or alternatively, the number of subtended facets on the surface of the cylinder. Furthermore, we can control the size and shape of the facets to achieve different goals, such as decrease computation, memory and storage requirements, or select a more beneficial cylinder discretization that allocates a greater number of facets in more interesting regions of the cylinder.

Table 5.2 – Summary of the parameters used to generate the virtual image in the simulations and in the experiments.

Parameter	Simulation	Experiments
Cylinder length [mm]		200
Cylinder radius [mm]		25
Camera distance [mm]		250
Focal length [mm]	30	24
Image resolution [px]	2000 x 2000	5616 x 3744
Number of facets ϕ		64
Number of facets θ		81
Total number of facets N		5184




Figure 5.14 – An image of the sample and cylinder, simulated with Mitsuba. (a) shows the rendered HDR image, which is perfectly aligned with the camera. We can also see the ϕ and θ axes of the spherical coordinate system attached to the center of the sample. ϕ spans the circumference, and θ the height of the cylinder. Because the coordinate system is attached to the sample, the parts of the cylinder closer to the sample have higher θ angles. (b) shows the facets of the virtual cylinder when projected onto the image. These facets are used to create the virtual image.

We start by generating a virtual cylinder, that has the same dimensions as the one used in the simulations and experiments. The cylinder parameters can be found in the top part of Table 5.2. We then discretize the cylinder into a desired number of facets along its length and its circumference. We use ϕ and θ to denote the facets along the cylinder circumference and length, respectively. We use a spherical coordinate system attached to the center of the sample, since the BTDF of the sample expressed in its spherical coordinates is ultimately what we are interested in measuring. A uniform discretization of the virtual cylinder with the parameters from Table 5.2 can be seen in Figure 5.9.

Given the cylinder discretization and the internal and external camera parameters, that we have from the geometric calibration, we can now project the facets onto the camera image. The correct alignment of the system components, that we do in the calibration step, is very important here. Figure 5.14a shows an example simulated image, and Figure 5.14b shows the facets of the virtual cylinder when projected onto the image.

Now we can create facet masks that we use for computing the virtual image pixels. For each image pixel, we check inside which projected facet does it find itself. As long as the measurement geometry (sample, cylinder and camera positions) does not change, we can reuse the same facet mask.




Figure 5.15 – Virtual image and cylinder unrolling. (a) The virtual image computed from the simulated image shown in Figure 5.14a. The values of the virtual pixels are the mean values of the image pixels for each facet in Figure 5.14b. (b) The process of unrolling the cylinder. This map represents the (inside) wall of the cylinder shown in Figure 5.9, where the opening at the sample side is in blue and the opening at the camera side is in yellow. Both figures share the same axes.

A virtual pixel's value is calculated as the mean value of all image pixels that belong to the corresponding facet. Figure 5.15a shows a virtual image computed from the image shown in Figure 5.14a with the facet masks of Figure 5.14b. Figure 5.15b shows the discretized cylinder from Figure 5.9 after it has been cut along its length, and then unrolled onto a flat surface. The color correspond between the two figures, blue represents the cylinder opening at the sample side, and yellow the cylinder opening at the camera side. There is a one-to-one mapping between Figures 5.15a and 5.15b, in the sense that the discretization of 5.15b, after projection onto the image (Figure 5.14b), created the masks that determine which pixels contribute to which virtual pixel's value, and the virtual pixel values are shown in 5.15a.

5.5.4 Initial facet irradiance from equilibrium radiance

The virtual image represents the steady-state or equilibrium radiance captured by the camera for the uniform facet discretization, instead of the original pixel discretization, of the cylinder. To measure the BTDF of the sample, we have to go from the steady-state radiance captured by the camera to the radiance emitted by the sample towards each facet. Alternatively, we can recover the BTDF from the irradiance of the facets that can be attributed solely to the sample radiance, because in equation 5.45, we showed that there is a linear relationship between the two.

Our approach for computing the facet irradiance due to the sample from the equilibrium radiance is similar to the approach in dual photography [89], which is based on the Helmholtz reciprocity of light. We already discussed the light transport matrix in Section 5.4. However,

instead of transposing the LTM, we invert it. Transposing the matrix will switch the positions of the sample and the camera and present us with an image as if it were "captured" by the sample and "illuminated" by the camera. Instead, we would like to "cancel" the light transport inside the cylinder, i.e., remove all light bounces except the first one. Therefore, we are looking for the initial facet irradiance that will produce the equilibrium radiance captured by the camera.

Given a Lambertian cylinder with a completely diffuse reflection, we can calculate the initial facet irradiance by inverting equation 5.48

$$\mathbf{e}_0 = \frac{\pi}{\rho} \left(\mathbf{I} - \frac{\rho}{\pi} \mathbf{G}^L \right) \mathbf{I}^{-\mathbf{p}_0}. \quad (5.56)$$

However, perfectly Lambertian materials do not exist. Therefore, we discussed a much more involved computation to produce the LTM for real materials. Nevertheless, to recover the initial facet irradiance, we apply the same principle, and invert the light transport matrix \mathbf{T} from equation 5.49:

$$\mathbf{e}_0 = \mathbf{T}^{-1} \mathbf{I}^{-\mathbf{p}_0}. \quad (5.57)$$

5.5.5 BTDF estimation

The final step is to compute the BTDF from the initial facet irradiance \mathbf{e}_0 . We start by expanding equation 5.45:

$$\begin{aligned} E_0(\mathbf{p}_i) &= L(\mathbf{p}_s \rightarrow \mathbf{p}_i) G(\mathbf{p}_s \leftrightarrow \mathbf{p}_i) A(\mathbf{p}_s) \\ &= E_i f_t(\mathbf{p}_l \rightarrow \mathbf{p}_s \rightarrow \mathbf{p}_i) G(\mathbf{p}_s \leftrightarrow \mathbf{p}_i) A(\mathbf{p}_s), \end{aligned}$$

where E_i is the collimated irradiance on the sample from the light source, \mathbf{p}_l is the center of the light source, and $f_t(\mathbf{p}_l \rightarrow \mathbf{p}_s \rightarrow \mathbf{p}_i)$ is the BTDF of the sample. The relations between the quantities from the equation are illustrated on Figure 5.16. From here we can easily compute the BTDF:

$$f_t(\mathbf{p}_l \rightarrow \mathbf{p}_s \rightarrow \mathbf{p}_i) = \frac{E_0(\mathbf{p}_i)}{E_i G(\mathbf{p}_s \leftrightarrow \mathbf{p}_i) A(\mathbf{p}_s)}, \quad (5.58)$$

or in matrix form we can express it as:




Figure 5.16 – Sample irradiance E_i due to radiance from a collimated light source. The sample transmits the incident irradiance towards the cylinder facets in the form of radiance $L_i(p_s \rightarrow p_i)$. The radiance distribution is governed by the BTDF $f_t(p_l \rightarrow p_s \rightarrow p_i)$.

$$\mathbf{f}_t = \frac{1}{E_i} \mathbf{e}_0 \oslash \mathbf{g}_s, \quad (5.59)$$

where \oslash is Hadamard (element-wise) division, and $\mathbf{g}_s = (G(p_s \leftrightarrow p_1) A(p_s), G(p_s \leftrightarrow p_2) A(p_s), \dots)^T$.

The irradiance on the sample E_i is measured during the radiometric calibration, by taking a photo of the collimated light source when there is no sample in place. Since the light source is collimated, the irradiance on the image sensor will be the same as the one on the sample. The camera measures only relative radiance values, and not absolute, as it is not radiometrically calibrated. However, the values of \mathbf{e}_0 are also scaled by the same factor, therefore, they are comparable to those of E_i . Naturally, all values should be scaled appropriately for the exposure period of the photo.

5.6 Simulations

In this section we will discuss our findings when evaluating our proposed system for BTDF estimation in simulations performed with Mitsuba: A Physically Based Renderer, Jakob [49]. Mitsuba is a research-oriented physically based rendering system. It is extremely modular, as it consists of a small set of core libraries and over 100 different plugins that implement functionality ranging from materials and light sources to complete rendering algorithms. It is a very convenient tool for us, as it offers great control over the light and material interactions.




Figure 5.17 – The physically based renderer Mitsuba allows us to accurately simulate real-world light interactions with objects. The three example images were produced by using different settings for the number of light bounces inside the cylinder, or the scene materials have different characteristics (BRDF, BTDF), while the scene geometry remains the same.

5.6.1 Simulations with Mitsuba: A Physically Based Renderer

We recreated our BTDF estimation setup from Figure 5.2 as a 3D scene in Mitsuba. Since Mitsuba is a physically-based renderer, it simulates the interaction of light with objects truthfully. All images in Figure 5.17 were rendered with Mitsuba by using different parameters and plugins.

The plugin developed by Dupuy and Jakob [22] allows us to load the measured BRDF of Spectralon, which can be found in the open-access material database⁴ of the Realistic Graphics Lab (RGL) at EPFL. We also measured the BTDF of our samples with a very accurate gonio-photometer, the pgII from Pab Advanced Technologies Ltd, to which Wenzel Jakob of the RGL kindly gave us access, and helped us tremendously with the measurements, to which we are grateful.

We modified Mitsuba to read a measured BTDF of our samples, and apply it as a light distribution of an area light source. We would then make the sample an area light source. Instead of this shortcut, we could have recreated the scene as if it would be in reality. In that case the sample would receive its proper BTDF, and it would be illuminated by a collimated light source. We, however, opted for the first approach, since it produces images with much less noise. That is the case because each shadow ray cast towards the light source after any number of bounces would reach the sample (the light source), since it is visible from any point on the cylinder. If, as in the second case, a collimated light source illuminated the sample, a ray must first reach the sample for the light source to become accessible. Like this, we would need to trace many more samples per pixel to reach the same level of signal-to-noise ratio (SNR).

⁴<http://rgl.epfl.ch/materials>




Figure 5.18 – The light directions ω are represented in the spherical coordinate system attached to the sample. The azimuthal angle ϕ is the angle from the x-axis to the projection of the ω direction on the xy-plane. The inclination, or polar angle θ is the angle between the z-axis and the direction ω . The hemisphere of transmitted directions \mathcal{H}_t^2 faces the camera, and the hemisphere of incident directions \mathcal{H}_i^2 (not shown) faces the light source, $\mathcal{H}_i^2 = -\mathcal{H}_t^2$.

5.6.2 Simulation parameters

The incoming and outgoing hemispheres of directions of radiance are measured in a spherical coordinate system centered on the sample. For all simulations and experiments we used an incident direction aligned with the axis of the cylinder, thus $\theta_i = 0^\circ$ and ϕ_i is irrelevant, since this is the pole of the hemisphere. The hemisphere of outgoing directions is rotated towards the inside of the cylinder, where θ_o spans the height of the cylinder, and ϕ_o its circumference. Figure 5.18 shows the spherical coordinate system attached to the sample for the outgoing directions \mathcal{H}_t^2 .

The B(R/T)DF can be isotropic or anisotropic. Most materials are isotropic. Their BRDF changes when the illuminating direction changes its θ_i angle, but it does not vary when it is illuminated from a different ϕ_i . Alternatively, we can rotate the material about its normal without changing its appearance. An example of isotropic BRDF is Spectralon. Anisotropic materials have a more general BRDF, since it changes with both the incident θ_i and ϕ_i angles. Anisotropic materials are less common, and some examples include brushed metals and textiles.




Figure 5.19 – Ground truth BTDF of (a) a completely diffusing (Lambertian) sample and (b) a IJM 617 sample, which is a coated white polyester textile. The diffuse sample is synthetic, and it has a constant BTDF. The BTDF of IJM 617 was measured with a pgII goniophotometer.

There is a special case where the B(R/T)DF can be constant. This is called a *Lambertian* B(R/T)DF, and we will also use the term *diffuse* interchangeably. The value of the Lambertian BRDF in every direction of the unit hemisphere is ρ/π , where ρ is the material reflectance, and π is the projected solid angle of the hemisphere (the area of a unit circle). A material that transmits light equally in all directions has a Lambertian BTDF. Therefore, its BTDF is constant in all directions of the unit hemisphere, with a value of T/π , where T is the transmittance of the material. It is unknown if such materials exist in the real world. In the simulations, we used $\rho = 0.99$ and $T = 1$. The constant BTDF will therefore be $1/\pi = 0.3183$.

The ground-truth BTDF of a synthetic Lambertian sample and a real IJM 617 sample can be seen in Figure 5.19. The horizontal axis shows ϕ_o and the vertical axis shows θ_o , in the sample's hemisphere of outgoing directions \mathcal{H}_t^2 . The Lambertian sample has a constant BTDF, since it perfectly diffuses the light in all directions. The IJM 617 sample is a coated white polyester textile used for printing. It diffuses light fairly well, but there is a visible falloff towards steeper outgoing θ_o angles.

Table 5.3 summarizes the parameters that we used for the simulations and for the experiments. We tried to keep them consistent where relevant, such that the results will be comparable.

5.6.3 Single bounce

In this case the light leaves the sample, reflects from the cylinder and is registered by the camera. This simplified case is interesting to analyze because the cylinder can be considered as an infinite set of adjacent lines. Therefore, the cylinder can be replaced by a long and flat stripe of diffusely reflecting material that rotates around the sample. There will be no interreflections, thus the single light bounce will greatly simplify the mathematical formulation for retrieving

Table 5.3 – This table summarizes the experimental setup, camera, rendering and virtual image parameters for the Mitsuba simulations and for the experiments.

Parameter	Simulation	Experiments
Cylinder length [mm]		200
Cylinder radius [mm]		25
Sample radius [mm]		3
Camera distance [mm]		250
Cylinder coating	Diff., Spec.	Spectralon
Focal length [mm]	30	24
Image resolution [px]	2000 x 2000	5616 x 3744
Aperture size	f/ ∞	f/2.8
Number of bounces	[1, ∞]	∞
Samples per pixel	4096	n/a
Number of facets ϕ		64
Number of facets θ		81
Total number of facets N		5184

the BTDF. This approach will require multiple photos of the stripe taken at different rotation angles ϕ_o . Each photo will contain many outgoing θ_o angles and a single outgoing ϕ_o angle.

In this case we can still use equation 5.57, however, we only need to invert \mathbf{T}_1 to compute \mathbf{e}_0 . \mathbf{T}_1 is a diagonal matrix, composed of the BRDF of each facet evaluated with the sample as an incident direction and the camera as an outgoing direction (equation 5.50).

The BTDF recovered from the rendered images that were allowed a single light bounce from the cylinder can be seen in Figure 5.20. The left column features the hypothetical diffuse (Lambertian) sample, and the right column a real sample from the material Canon IJM 617. The top row shows the ground truth BTDF, the middle row shows the BTDF recovered from the images simulated using a Lambertian cylinder, and the bottom row shows the BTDF recovered from images simulated using the Spectralon cylinder.

When using a Lambertian cylinder, we can see that the estimated BTDF is very accurate for both samples. The mean difference from the ground-truth is 0.09 % and 0.15 % for the diffuse and the IJM 617 samples, respectively. We can also see that the directional variation of the IJM 617 sample is accurately recovered, in a slight smoother version than the measured ground-truth. The smoothing is due to the non-infinitesimal size of the sample; its finite size illuminates several facets even when considering a single collimated outgoing light direction.

In the case of the Spectralon coated cylinder, the accuracy decreases somewhat. The mean difference to the GT is now 0.35 % and 0.44 %, for the diffuse and IJM 617 samples, respectively. We can see stripes in the BTDF of the diffuse sample, when the GT is completely uniform in all directions. This is likely caused by measurement errors (e.g., noise) in the BRDF of Spectralon, compounded by the size of the facets. We can also see the errors in the BTDF of IJM 617. While




Figure 5.20 – Recovered BTDF after a single light bounce inside the cylinder. The left column features the hypothetical diffuse (Lambertian) sample, and the right column the real sample IJM 617. The top row shows the ground truth BTDF, the middle row shows the BTDF recovered from the simulated images when using a Lambertian cylinder, and the bottom row shows the recovered BTDF from images simulated using the Spectralon cylinder.




Figure 5.21 – Recovered BTDF after a single light bounce inside the cylinder. These BTDFs feature high-frequency patterns with abrupt radiance changes. The left column features a BTDF in the shape of an arrow pointing north-east, and the right column a BTDF in the shape of the EPFL logo. The top row shows the ground truth BTDF, the middle row shows the BTDF recovered from the simulated images when using a Lambertian cylinder, and the bottom row shows the recovered BTDF from images simulated using the Spectralon cylinder.

the shape of the BTDF is generally correct, it is slightly distorted where the stripes appear on the diffuse sample.

Figure 5.21 shows the estimated BTDF of two synthetic samples that feature high-frequency patterns with abrupt changes. There are no obvious differences between the reconstructions with both the Lambertian and the Spectralon cylinders. The Arrow sample spans almost all outgoing angles that the cylinder covers. We can see that the reconstruction in the lower part of the figure, larger θ_o angles, is sharper and more accurate. These outgoing angles send light in the region of the cylinder that is closer to the sample. This outgoing radiance is intercepted by the side of the cylinder close to a right angle, and it therefore covers less facets. When θ_o decreases, the angle that the light forms with the cylinder sides is much steeper, and thus the projection of the light beam covers more facets, according to Lambert's cosine law. We see a similar situation with the EPFL logo, where the top of the letters get stretched, and this smoothing decreases the contrast of the recovered BTDF.

5.6.4 Infinite bounces

In Figure 5.22 we can see the estimated BTDF after an infinite bounces of light inside the cylinder. This case should simulate what happens in the real world. However, even in reality light does not reflect an infinite number of times before it is absorbed or it escapes the cylinder, but a few tens of times, depending on the geometry and absorptance of the material. With each absorption and scattering event, the initial radiance leaving the sample gets diminished, and therefore contributes less to the overall radiance leaving towards the camera.

Diffuse sample, Lambertian cylinder is the simplest case. Due to the symmetry of our scene, for the diffuse sample it sufficed to render a single slice of the image (cylinder), and then repeated it circularly to create the complete image. This allows us to trace 2^{20} samples per pixel, while the rendering still finishes within a few minutes. Like this, we effectively render only 1000 pixels, instead of 4 million pixels, for a 2000 x 2000 pixel image.

The mean difference to the ground truth (GT) in this case is 3.33 %. As we can see from the figure, the error is mostly concentrated in the region close to the camera, up until a θ_o of about 11° . The facets closer to the camera are more dependent on light interreflection (number of bounces) to reach their equilibrium radiance, since they receive less direct light from the sample. Thus, it is likely that neighboring pixels reach different numbers of bounces, which creates a significant amount of noise in their values.

If we look at the estimated BTDF for the diffuse sample in the Spectralon coated cylinder, we can see that the standard deviation of the values decreases, and the mean difference is now also lower at 1 %, compared to the same sample in the Lambertian cylinder. We can still see a similar gradual decrease of the BTDF intensity towards lower θ_o angles. If the reflectance of the material would decrease further, this BTDF intensity decrease would be eliminated, since the light will achieve a lower number of bounces.




Figure 5.22 – Recovered BTDF after an infinite light bounces inside the cylinder. The left column features the hypothetical diffuse (Lambertian) sample, and the right column the real sample IJM 617. The top row shows the ground truth BTDF, the middle row shows the BTDF recovered from the simulated images when using a Lambertian cylinder, and the bottom row shows the recovered BTDF from images simulated using the Spectralon cylinder.




Figure 5.23 – Recovered BTDF after an infinite light bounces from the cylinder. These BTDFs feature high-frequency patterns with abrupt intensity changes. The left column features a BTDF in the shape of an arrow pointing north-east, and the right column a BTDF in the EPFL logo. The top row shows the ground truth BTDF, the middle row shows the BTDF recovered from the simulated images when using a Lambertian cylinder, and the bottom row shows the recovered BTDF from images simulated using the Spectralon cylinder.

A more prominent artifact in this case is the bulging in the center of the BTDF, 10° to 25° θ_0 . What this tells us is that, according to our method, it takes larger values for the BTDF in those regions to produce the image values of the rendered image. Those are also the cylinder regions that have the highest likelihood to reach the greatest number of light bounces. We do not get the bulging with the Lambertian cylinder, but we calculate the light transport matrix differently, as explained in subsection 5.4.1. As we get closer to the sample, the BTDF reaches its correct value. In those regions the direct illumination from the sample dominates the light due to interreflections.

When estimating the IJM 617 sample's BTDF with the Lambertian cylinder, we can see that the subtle details of the BTDF are lost. The mean difference from the GT is 1.76 %. The BTDF is again estimated to have a slightly lower average value, similar to what we saw with the diffuse sample. The general shape of the BTDF is, however, preserved. A Lambertian cylinder with a reflectance value of 75 %, instead of 99 %, yields a BTDF which preserves the finer details. This, however, lowers the radiance captured by the camera, which would require more sensitive sensors, or longer exposure times to keep the same SNR.

The final subfigure shows the reconstruction of the IJM 617 sample's BTDF with a Spectralon coated cylinder. The subtle details again seem to be lost. In this reconstruction, we can still see a shape that more closely approaches the ground truth BTDF of the sample, with the values in the top-left and top-right corners being slightly darker, than the top-center part of the figure, similar to the ground truth BTDF. The mean difference from the ground truth is 0.59 %. We again see the characteristic bulge in the middle of the cylinder, that we saw when estimating the BTDF of the diffuse sample with the Spectralon coated cylinder.

In Figure 5.23 we can see the reconstruction of the synthetic and difficult Arrow and EPFL logo samples. The reconstructed BTDFs again seem correct, as in the single bounce case, without major differences. The most obvious one is the increase in the noise level at the top of the letters of the EPFL logo in the case of the Lambertian cylinder. This is again inline with what we saw in the BTDF reconstruction of the IJM 617 sample.

5.6.5 Compensating for the systematic error

Figure 5.24 shows the BTDF reconstruction results after applying a simple compensation for the systematic error that causes the bulging of the BTDF in the central part of the cylinder. To compute the compensation we perform forward ray tracing by using equation 5.49. We use the BTDF of the *diffuse* sample to compute the initial facet irradiance \mathbf{e}_0 and the light transport matrix \mathbf{T} for the Spectralon cylinder. This computes the virtual image for the diffuse sample. The compensation is then a simple ratio between the virtual image that we compute and the virtual image rendered by Mitsuba. We then apply this compensation to the virtual image rendered by Mitsuba for *any* sample. Using the compensated virtual image, we compute the BTDF with the inverse approach, i.e., by following the steps backwards.




Figure 5.24 – Recovered BTDF after an infinite light bounces inside the cylinder, and after a compensation has been applied for the systematic error of the central BTDF values. The left column features the hypothetical diffuse (Lambertian) sample, and the right column the real sample IJM 617. The top row shows the ground truth BTDF, and the middle row shows the BTDF recovered from the simulated images when using a Spectralon cylinder. The bottom row shows the recovered BTDF from images simulated also using the Spectralon cylinder, but to which the simple compensation has been applied, and are lightly filtered with a Gaussian filter (size = 5, $\sigma = 2$) to eliminate the noise.

In the bottom row of Figure 5.24 one can see the BTDF which we compute after applying a compensation to the virtual images computed by Mitsuba. The compensation eliminates all errors for the diffuse sample, which is to be expected, since we use it as the basis for computing the compensation. What is more interesting is the BTDF of the IJM 617 sample, to whose virtual image rendered with Mitsuba we applied the compensation computed with the diffuse sample. We also lightly filtered the result with a Gaussian filter of size = 5 and $\sigma = 2$ to eliminate the noise. We can see that the simple compensation works well, and that even the subtle directional variations of the BTDF have reappeared. Results for four additional samples, Canon IJM 601, Canon IJM 637, Avery Dennison MPI 2020, and Avery Dennison MPI 3151, can be found in Appendix 5.B.

5.7 Experiments

In this section we discuss the experimental setup that we used to evaluate the proposed radial imaging approach for estimating the BTDF of real samples. The experiments require more steps in order to process the images and compute the BTDF, than the simulations. The added complexity comes from the preprocessing steps, that are imposed by limitations in the hardware, and the difficulty of aligning the optical elements.

5.7.1 Experimental setup

For the experiments, we recreated the setup that can be seen in Figure 5.2. All of the parts that we used were off-the-shelf optical components, except for the cylinder made out of Spectralon, which we acquired fairly easily from SphereOptics. We performed our experiments in a dark room to minimize the amount of stray light. The dark room can be easily replaced by a light insulating box with an opening that provides access to the measured samples.

An image that contains all components which comprise our experimental setup can be seen in Figure 5.25. The components before the cylinder are used as a simple collimator. As a light source we use a Dolan-Jenner DC-950H DC-Regulated Halogen Illuminator to which we attach an optical fiber. The output of the optical fiber goes through an iris diaphragm, then gets filtered by an interference filter with a center wavelength of 545 nm and a full width at half maximum (FWHM) of 10 nm. A 100 mm lens has the iris diaphragm at its focus, and sends the light in parallel rays through another iris diaphragm that we use to control the radius of the illumination spot on the measured sample.

The sample is mounted on a black disk with an opening at the center. The disk is mounted at the opening of the cylinder, and the sample sits flush with the beginning of the cylinder. The light that passes through the sample gets scattered inside the cylinder, and the camera, Canon 5D Mark II with a Canon 24 mm f/2.8 lens, captures the steady state radiance.




Figure 5.25 – An implementation of our proposed radial imaging setup. The elements to the right of the cylinder are used to collimate the illumination light beam. The light that is transmitted by the sample is collected by the cylinder. After a large number of interreflection, the cylinder radiance reaches a steady-state, which is captured by the camera.

5.7.2 Image capture

The photos captured by the camera have a much lower dynamic range than the range of radiances inside the cylinder. Therefore, we have to capture multiple photos with different exposure times and merge them in a single HDR image. Table 5.4 shows the camera parameters that we used to capture a set of images that will be merged into an HDR image.

Before we assemble the captured photos in an HDR image, we perform vignetting correction to each photo. We also geometrically calibrate the camera, because we need the internal camera matrix to project the virtual facets onto the image, and to undo the radial and tangential distortions of the image caused by the lens. These steps are described in more detail in subsection 5.5.2.

Table 5.4 – Summary of the camera parameters used for capturing photos for the experiments.

Parameter	Value
Image type	Linear RAW
Aperture size	f/2.8
Exposure time [s]	$\frac{1}{6}$, $\frac{1}{4}$, 0.5, 1, 2, 4, 8, 15, 30
ISO	100




Figure 5.26 – The alignment of the cylinder to the camera is not perfect, which is why we had to lightly manipulate the virtual cylinder before projection onto the image. (a) An HDR image computed from a set of images captured with different exposure times. (b) A virtual cylinder projected onto the HDR image, after a small rotation and translation.

5.7.3 Results

In Figure 5.26a we can see the HDR image computed from a set of images captured with the camera parameters from Table 5.4. In Figure 5.26b we can see the projection of the virtual cylinder onto the image. Even after the calibration, the alignment of the sample, cylinder and camera is not perfect, as it was in the Mitsuba simulations. We therefore performed a manual rotation and translation of the virtual cylinder before projecting it onto the image. The rotation angles are small, 0.29° and 0.75° along the x- and y-axes, respectively, and the translation along the x-axis is 0.075 mm. The camera was also translated by around 5 mm further away from the sample. If the sample, cylinder and camera are not aligned properly, and at the correct locations and orientations that were used when computing the LTM, then it becomes invalid. Therefore, a precise alignment of all optical elements is crucial.

Figure 5.27 shows the comparison between the virtual image computed from the captured HDR image (blue line), and the virtual image computed from the image simulated with Mitsuba (red line), using the same scene parameters. Here we can show them as lines, instead of surfaces of two variables, since the lines represent the average radiance across all ϕ angles for each θ angle. The blue lines have been rescaled such that their maximal values are equal to those of the red lines. Naturally, the scale of the y-axis is for the simulated image (red line).

The alignment on the right side of the plots (close to the sample) is good across all samples. However, as we move towards the left side (closer to the camera), they start to diverge. The largest differences appear in the middle of the plots, where the values of the captured images are higher than if the simulated images. Finally, as get to the left side, the values converge again. The only exception from this observation is the IJM 617 sample, and it is also the only textile sample.




Figure 5.27 – A comparison between the captured HDR images (shown with blue line) and the images simulated with Mitsuba (shown with red line) using the same parameters. The lines represent the average radiance across all ϕ angles for each θ angle. The left sides of the plots are close to the camera and the right sides are close to the sample. We can see a similar trend for most samples, where the ends are well aligned between the two images, but the captured images have higher radiance in the central part of the cylinder. A notable exception is IJM 617, where the differences are also pronounced on the left side (close to the camera).




Figure 5.28 – Required compensation that will equalize the captured and the simulated images. There is a systematic error in the measurements, with the IJM 617 being the only outlier, and also the only textile.

The sharp falloff at the left end of the plots (the camera end of the cylinder) for the captured images is due to the finite size of the physical aperture in the camera. The camera is focused on the sample, and the cylinder end close to the camera gets blurred. The blur circles then mix the content of the image due to parts of the cylinder and parts outside of the cylinder. This problem can be addressed by decreasing the size of the aperture, and by taking multiple HDR images with different focal planes, e.g., at the sample, at the middle of the cylinder, and at the cylinder opening close to the camera.

To confirm the observed trend, we computed the ratio between the two lines, for each sample. The results are shown in Figure 5.28. These values are, as in the previous figure, the average across all ϕ angles, for each θ angle. This is the required compensation by which we need to divide the captured values to reach the simulated values. The figure confirms that there are indeed systematic differences in the measurements. With the exception of IJM 617, the samples' virtual images could be compensated with a very similar factor.

There are multiple factors that can contribute to these systematic differences. As it can be seen in Figure 5.26a, the sample size in the experiments is larger than the one in the simulations. Although the sample is illuminated with a collimated beam of light, of the same size as in the simulations, the increased size of the sample contributes to backscattering of light inside the cylinder. In the simulations, this light would escape the cylinder. Although it does not seem to influence the image close to the sample, it does influence the values closer to the camera, since those are much more dependent on the interreflections, than on direct irradiance by the sample.

Furthermore, subsurface scattering inside the sample increases its effective illuminated area. Therefore, the sample must be cut to the correct size, or be isolated from backscattering the light inside the cylinder by covering it with a physical barrier or highly absorbing paint.

Another source of backscattering that does not appear in the simulations is from the camera lens towards the cylinder. This can also contribute towards increasing the radiance values, especially of the pixels near the camera. The camera should thus either be mounted further away from the cylinder and/or decreased in size.

Spectralon is translucent, which can be seen in Figure 5.30b. Therefore, light that enters the material at one point, can get transported under the surface, to emerge at a different point. This can also add to the differences in the measurements, since the simulations were carried out without subsurface scattering. This would mean that the high intensity irradiance close to the sample could be transported towards the middle of the cylinder, where it emerges and increases the registered radiance by the camera.

Another reason that might add to the discrepancies in values between the captured and the simulated images is the difference in surface roughness of the Spectralon that was used to make the cylinder and that of the measured Spectralon tile. The Spectralon BRDF measurement that we used in the Mistuba simulation might have been performed on a tile that was compacted more than the Spectralon of which the cylinder was made.

Figure 5.29 shows the ratio between the BTDF computed from the captured virtual images, and the ground truth BTDF of the materials, measured with the pab ppl goniophotometer. Before computing the BTDF, we compensated each virtual image by the corresponding amount, shown in Figure 5.27. Note that this compensation is computed as the average over all ϕ angles, and not for every (ϕ, θ) pair.

We can again see a systematic error in the results. They are lower close to the sample, falling below 10 %, and increase to more than 300 % close to the camera. The main reason is the misalignment issue between the sample, cylinder and camera, which we discussed above. The light source seems to illuminate the samples slightly off-center, shifted along the x-axis, as it can be seen by the increased intensity around $\phi = 190^\circ$, and decreased at the opposite side of the cylinder $\phi = 10^\circ$. This shows that the approach is highly sensitive to misalignment issues.




Figure 5.29 – The BTDF ratio when it is computed from the captured images, to the ground truth. Each captured image was compensated by the corresponding amount, as shown in Figure 5.28. The differences are below 10 % in the parts close to the sample, and increase to more than 300 % when moving towards the camera. We again see a systematic error in all figures, which is caused by the misalignment of the sample, cylinder and camera in our experimental setup.

5.8 Practical considerations and limitations


Generating the light transport matrix is very computationally and memory intensive. On the positive side, all operations are additions and multiplications. The graphics processing unit (GPU) is a specialized hardware to very efficiently perform such operations. The problem is also embarrassingly parallel, and it can easily be run on multiple GPUs, as long as there is enough memory. Selecting a symmetric scene, like ours, means that we need to compute the columns of the LTM corresponding to just a single ϕ_o angle and all θ_o angles of the cylinder, i.e., a stripe of facets along its height. The rest of the matrix can be populated with circular shifting of this data.

We computed the light transport matrix for 5184 facets and 100 light bounces in 7.14 hours on three Nvidia Titan Xp GPUs, running optimized MATLAB code. This can certainly be accelerated with better optimization in a lower-level coding language. The main difficulty with the computation is storing the 3D BRDF tensor of the cylinder (Figure 5.11). Again, due to the symmetry of the scene, it can be stored partially, and the rest can be populated with circular shifting of this data. This is acceptable, since the light transport matrix is computed offline, and only once as long as the scene geometry and cylinder BRDF do not change. The forward rendering, computing the equilibrium radiance from the BTDF, is a simple matrix-vector product (equation 5.49), and the inverse rendering, computing the BTDF from the equilibrium radiance, is a matrix inversion followed by a matrix-vector product (equation 5.57). Since the light transport matrix is fixed, the inverted matrix can also be stored, thus the inverse rendering will also become a simple matrix-vector product.

The proposed approach is capable of measuring many outgoing angles, but only a single wavelength of light. For each additional wavelengths, we will need to exchange the light filter, and repeat the process of taking a set of photos. With the parameters from Table 5.4, we were able to take one set of photos in around one minute. A more practical approach to illuminate with white light, and use a multispectral or a hyperspectral camera to capture all desired wavelengths in a single capture session. This will trade-off spatial for spectral resolution. Given the size of the virtual pixels, multiple image pixels will still be available for averaging to compute the value of the virtual pixel.

The size of the sample acts as a low pass filter on the computed BTDF (Figure 5.21). That makes it difficult to capture intricate details in the BTDF. Decreasing the radius of the sample can preserve more details, however, the cylinder will then become darker. A more sensitive camera, or longer exposure times must then be used for achieving decent SNR.

In the current implementation, we illuminate the sample at a right angle ($\theta_i = 0^\circ$). For samples that are fairly transparent, there can be problems with lens flare, since most of the light will continue straight towards the camera, and the smaller portion will be scattered around the cylinder. The stray light in the camera optics will make the measurements biased and therefore useless. This is illustrated in Figure 5.30a. To avoid this problem, a mirror tilted at 45° can be placed in front of the camera to divert the direct light component away from the camera, and



(a) Lens flare in the camera optics due to the unscattered portion of the direct illumination

(b) Illumination through the cylinder wall with a phone flashlight

Figure 5.30 – Practical limitations of our proposed setup for estimating BTDF. (a) When measuring samples that are highly transmitting and lightly scattering, the intense unscattered portion of the illumination causes lens flare in the camera optics. This stray light scatters to other pixels of the image, and biases the measurement. It is most notable as the halo around the central saturated part of the image, when this portion of the image should be dark. (b) The translucency of the Spectralon cylinder lets intense light enter from the outside. This imposes an upper limit on the intensity of the light source, for light to be fully contained inside the cylinder.

outside of the cylinder. This will also lower the captured radiance of the cylinder when using larger apertures, because a portion of the camera lens will be blocked.

Another practical consideration is that Spectralon is translucent. Thus, light could escape the cylinder if it is not thick enough for the irradiance exerted on its walls. Figure 5.30b depicts the opposite case, where light enters the cylinder from the outside. This upper limit on the intensity of the illumination could be circumvented by using more sensitive cameras.

5.9 Summary

BSDFs have traditionally been measured with scanning goniophotometers. Where measurement accuracy is critical, they are irreplaceable. However, they require long acquisition times. This issue lead to the design and development of parallel goniophotometers. Many different approaches have already been proposed, and new approaches are continuously being added, or the existing ones are being perfected.

In this work, we proposed and evaluated a novel design for a parallel goniophotometer, which is simple and inexpensive to build. Our approach relies on a cylinder with a highly diffusing internal surface, that reflects many outgoing directions from the sample towards the camera. Because the cylinder inside is convex, many interreflections take place before it reaches its steady-state radiance. We then invert the light transport that takes place inside the cylinder to compute the irradiance of the cylinder facets that is due only to the sample, and not due to other cylinder facets. From this initial irradiance we can compute the BTDF of the sample.

Our approach relies on a simple hardware design and an involved computational framework. We have therefore presented all steps from the mathematical foundation diligently and transparently. The intensive computations are performed offline, and can be reused for as long as the geometry of the measurement setup remains unchanged. The online computations performed during measurement are lightweight, consisting of simple matrix-vector products.

The results from the simulations look promising. Inverting the light bounces with the light transport matrix works well, and the computed BTDFs are well aligned with the ground-truth, for both low-frequency and high-frequency patterns with abrupt changes. We also identified some possible issues. The BTDFs becomes somewhat noisy in the regions of the cylinder that are far away from the sample. This happens when we use very high reflectance values for the cylinder walls, and can mostly be attributed to the noise in the rendering. Decreasing the reflectance decreases the noise due to interreflections, but also decreases the equilibrium radiance inside the cylinder. The systematic error in the central part of the cylinder remains to be addressed at a more profound level than the simple compensation that we proposed, although it proved to be very effective in suppressing them. The size of the measured sample presents another trade-off. Where a larger sample acts as a smoothing filter on the BTDF, especially farther away from the sample, it allows for higher internal cylinder steady-state radiance.

The results from the real world experiments are not entirely satisfactory. That is mainly because of shortcomings in the assembly and calibration of our setup, where we failed to align all optical elements properly. We believe that an industrial grade implementation can approach the accuracy that we achieved in the simulations. We have also identified some of the practical limitations, and have proposed possible directions for their mitigation.

5.A Appendix: Shape of \mathbf{F}_r

The structure of the vectors and matrices from equation 5.35:

$$\begin{aligned}\mathbf{l} &= \mathbf{l}_e + \mathbf{F}_r \mathbf{G} \mathbf{l} \\ &= \mathbf{l}_e + \mathbf{F}_r (\mathbf{g} \odot \mathbf{l}),\end{aligned}$$

$$\mathbf{l} = \begin{bmatrix} L(p_1 \rightarrow p_1) \\ L(p_1 \rightarrow p_2) \\ \vdots \\ L(p_1 \rightarrow p_N) \\ L(p_2 \rightarrow p_1) \\ L(p_2 \rightarrow p_2) \\ \vdots \\ L(p_N \rightarrow p_N) \end{bmatrix}, \quad \mathbf{l}_e = \begin{bmatrix} L_e(p_1 \rightarrow p_1) \\ L_e(p_1 \rightarrow p_2) \\ \vdots \\ L_e(p_1 \rightarrow p_N) \\ L_e(p_2 \rightarrow p_1) \\ L_e(p_2 \rightarrow p_2) \\ \vdots \\ L_e(p_N \rightarrow p_N) \end{bmatrix}, \quad \mathbf{g} = \begin{bmatrix} G(p_1 \leftrightarrow p_1) A(p_1) \\ G(p_1 \leftrightarrow p_2) A(p_1) \\ \vdots \\ G(p_1 \leftrightarrow p_N) A(p_1) \\ G(p_2 \leftrightarrow p_1) A(p_2) \\ G(p_2 \leftrightarrow p_2) A(p_2) \\ \vdots \\ G(p_N \leftrightarrow p_N) A(p_N) \end{bmatrix}, \quad \mathbf{G} = \begin{bmatrix} \mathbf{g}^T \\ \mathbf{g}^T \\ \vdots \\ \mathbf{g}^T \\ \mathbf{g}^T \\ \mathbf{g}^T \\ \vdots \\ \mathbf{g}^T \end{bmatrix},$$

$$\mathbf{F}_r = \begin{bmatrix} f_r(p_1, p_1, p_1) & 0 \dots & \dots 0 & f_r(p_2, p_1, p_1) & 0 \dots & \dots 0 & f_r(p_N, p_1, p_1) & 0 \dots & \dots 0 \\ 0 \dots & f_r(p_1, p_1, p_2) & 0 \dots & \dots 0 & f_r(p_2, p_1, p_2) & 0 \dots & \dots 0 & f_r(p_N, p_1, p_2) & \dots 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 \dots & \dots 0 & f_r(p_1, p_N, p_N) & 0 \dots & \dots 0 & f_r(p_2, p_N, p_N) & 0 \dots & \dots 0 & f_r(p_N, p_N, p_N) \end{bmatrix}$$

$$\mathbf{l} \in \mathbb{R}^{N^2}, \quad \mathbf{l}_e \in \mathbb{R}^{N^2}, \quad \mathbf{g} \in \mathbb{R}^{N^2}, \quad \mathbf{G} \in \mathbb{R}^{N^2 \times N^2}, \quad \mathbf{F}_r \in \mathbb{R}^{N^2 \times N^2}.$$

5.B Appendix: Additional results




Figure 5.31 – Recovered BTDF after an infinite light bounces inside the cylinder, and after a compensation has been applied for the systematic error of the central BTDF values. The left column features a sample from Canon IJM 601, and the right column a sample from Canon IJM 637. The top row shows the ground truth BTDF, and the middle row shows the BTDF recovered from the simulated images when using a Spectralon cylinder. The bottom row shows the recovered BTDF from images simulated also using the Spectralon cylinder, but to which the simple compensation has been applied, and are lightly filtered with a Gaussian filter (size = 5, σ = 2) to eliminate the noise.




Figure 5.32 – Recovered BTDF after an infinite light bounces inside the cylinder, and after a compensation has been applied for the systematic error of the central BTDF values. The left column features a sample from Avery Dennison MPI 2020, and the right column a sample from Avery Dennison MPI 3151. The top row shows the ground truth BTDF, and the middle row shows the BTDF recovered from the simulated images when using a Spectralon cylinder. The bottom row shows the recovered BTDF from images simulated also using the Spectralon cylinder, but to which the simple compensation has been applied, and are lightly filtered with a Gaussian filter (size = 5, $\sigma = 2$) to eliminate the noise.

6 Conclusion

6.1 Thesis summary

In this thesis we looked into the measurement of material properties, primarily with cameras. In the first part of the thesis we were concerned with material appearance in reflectance mode, i.e., when the illuminant and the observer are on the same-side of the material. A very important feature that governs the appearance is surface roughness. Therefore, in Chapters 2 and 3 we explored the use of structured light systems for measuring the surface structure of materials, which can then be summarized by the roughness measure. In the second part of the thesis, we were interested in measuring the transmitting properties of materials. In Chapter 4 we focused on the hemispherical transmittance, and the proper measurement of total transmittance. We then went a step further in Chapter 5, and proposed a novel device for the measurement of directional transmittance, which is represented by the BTDF.

In Chapter 2, we discussed the importance of calibration for stereo systems based on structured light projection. In addition to geometric calibration, which is essential, structured light systems that employ codes with continuous variation require a radiometric calibration, when there is a non-linear processing in the loop. It is generally the case that the projectors add the non-linear processing, since modern cameras can output linear images. Due to the popularity and benefits of codes with continuous variation, radiometric calibration is slowly becoming another essential part of the process.

The radiometric calibration increases the complexity of the process in terms of additional capture time and manual intervention, for example, exchanging the calibration board used for geometric calibration with one for radiometric calibration. Therefore, we proposed an algorithm that performs a geometric calibration of the projector-camera pair and a radiometric calibration of the projector, simultaneously, using a single calibration board and from the same set of captured images. The algorithm does not require additional user intervention for the radiometric calibration. We also showed that our algorithm offers high quality calibration, similar to the state-of-the-art algorithms, for both calibrations.

Chapter 6. Conclusion

In Chapter 3 we applied the calibration procedure from Chapter 2 to a structured light system that we built for the purpose of measuring surface roughness. Macroscopic surface roughness is an important factor in determining the appearance of a surface. We focused on scanning the surface of printing textiles. Due to the increasing popularity of digital printing on textiles, it would be of great interest to devise color prediction models that can incorporate such surface information.

In the chapter we provided an overview of the main categories of structured light encoding techniques, and discussed the potential sources of errors. We then designed and assembled the system for depth estimation. We selected the surface roughness metric to compare the scans of our system to those from a profilometer based on the confocal principle. When properly calibrated, we managed to achieve comparable results for materials that had surface features with periods from 0.5 mm to 1 mm, surface height of up to 150 μm , and moderately complex surfaces. Moreover, our system offered faster scanning and/or wider scanning area than the profilometer, at the expense of lower resolution.

In the second part of the thesis we concentrated on measuring transmittance. In Chapter 4 we evaluated different techniques for measuring hemispherical transmittance. We presented different measurement geometries for total transmittance, and discussed the sources of error. We then built a system for total transmittance measurement to evaluate several of the measurement geometries. We also compared our measurements to those from commercial solutions. We showed that certain measurement geometries introduce a significant bias, which should be compensated. Therefore, when reporting total transmittance, it is a good practice to also state the measurement geometry.

Finally, in Chapter 5 we expanded the work on transmittance measurement, towards measuring the Bidirectional Transmittance Distribution Function (BTDF). It describes the angular distribution of the incident light after it passes through the material. It is thus of great interest for computer simulations, special effects and architectural modelling of illumination of spaces.

Our contribution is a novel design for a parallel goniophotometer. Our system is based around a diffusing cylinder that collects the outgoing light from a sample that is illuminated by a collimated beam of light. The steady-state radiance of the cylinder's interior is then captured by a camera. We then compute the BTDF after mathematically reversing all light interreflections between the cylinder walls. Our method trades mechanical complexity with computation. We provided a rigorous mathematical discussion on the foundation of the method. We then built the proposed design in Mitsuba, a physically based renderer, and truthfully simulated the physical interactions of light with the objects, to evaluate our approach. The results were well aligned with the ground-truth, opening the door for experimental validation. The prototype that we built was based on off-the-shelf components, except for the cylinder that can also be acquired easily. The experimental results were not entirely successful mainly due to mechanical alignment issues. However, we believe that an industrial grade implementation could overcome the practical issues, and approach the accuracy from the simulations.

6.2 Future research

6.2.1 Simultaneous calibration

Incremental calibration. The method does not mandate the calibration to be performed after a predetermined number of images have been acquired. Therefore, an incremental calibration can provide the operator with an immediate feedback on the calibration accuracy for the geometric calibration, and an overview of the gamma estimation convergence. This will allow to condition the number of captures on the desired calibration accuracy.

Self-identifying corners. Currently we expect to see all printed and projected corners in order to determine the orientation of the patterns and match the two sets of corners for computing the prewarp. However, an interesting future direction is to investigate pattern designs that can uniquely identify the corners by inspecting their surrounding squares.

Radiometric calibration for all color channels. In the current implementation, the radiometric calibration is performed only for one color channel, and it is assumed that it is equal for the other color channels. For devices that have widely varying response functions for different color channels, it is possible to generate projected pattern in different colors, and to extract their corner locations and intensities from the corresponding camera channels, and the white squares of the printed pattern. This will require a more intricate managing of possible failures when the projected pattern is not detected, and we also have to make sure that we do not interfere with the printed pattern.

Accounting for distance falloff to get more accurate values for the radiometric calibration. When the calibration board is relatively close to the projector and the camera, we should account for the irradiance changes due to the distance falloff. In this case, different parts of the projected pattern would lie at significantly different distances from the devices. Assume that the calibration board is close to a Lambertian reflector, we should normalize the captured values in the camera image by the distance squared from both devices. The geometric calibration provides us with the relative position and orientation of the calibration board with respect to both the projector and the camera, which is all the information needed to perform this compensation.

6.2.2 Surface roughness

The unexpected very narrow valleys that were visible in the reconstruction of the Canon IJM618, we suspect appeared due to specular reflections. The unpolarized light from the DLP projector gets partially polarized parallel to the surface when it reflects from the textile (a dielectric). Therefore, an effective solution for suppressing specular reflections is by using a linear polarization filter in front of the camera [70]. This, however, comes at the cost of a 4-8 times reduction in the captured brightness. Cross-polarization is an even more effective solution for lowering the amount of specularly reflected light and it uses polarization filters in front of both the camera and the projector [104].




Figure 6.1 – Geometry for capturing a different outgoing hemisphere.

Recent color prediction models for printing account for the roughness of the inks, but do not take into account the roughness of the substrate, which is about two orders of magnitude greater for textile substrates. Therefore, it would be interesting to extend the existing color prediction models to take into account the roughness of the surface at the macroscopic scale, by possibly incorporating local interreflections. Work in this direction has already been performed, for the more specialized, but very important case of specular V-cavities [86, 84].

6.2.3 Parallel goniophotometer

A simple extension to the proposed parallel goniophotometer design would be to explore different facet discretizations of the cylinder, since there are no inherent limits on the possible discretizations. Those with rotational symmetry are more efficient to compute, since the matrix can be computed only partially, and then circularly shifted and repeated. Interesting discretizations can allocate a larger amount of facets in more interesting regions of the cylinder. For example, we might be more interested in the smaller θ_o angles, and the current uniform discretization would work well in that case. A simple variation would be a discretization where the θ_o angles progress uniformly, which will make the facets closer to the camera larger.

There is, however, a limitation on the number of facets, because it can make the computation impractically long. A recent paper on differential rendering by Nimier-David et al. [75], could allow us to use much higher number of facets, and potentially the whole resolution of the camera. The concept is similar to ours, in the sense that it also tries to reverse the light transport in the scene. It starts from an initial state of a scene, and the goal is to update the parameters of the rendering such that it matches a target rendering. The downside is that the intense computation will have to be performed online, for each measured sample, but a static scene like ours could be set to an initial state that is not too far from the target state. Also, the differentiation requires the BTDF to be a smooth function, which should not be problematic for most materials.

As it can be seen from the results in Section 5.6, the design based on a cylinder covers all ϕ_o angles, however, the θ_o angles start from around 7° , and go all the way to more than 87° . The missing angles between 0° and 7° can be captured by rotating the light source and the sample in unison, while the cylinder and camera remain fixed. This is illustrated in Figure 6.1. Furthermore, to capture an additional BTDF slice for a different incident direction, the




Figure 6.2 – Geometry for capturing an oblique incoming direction.

light source should be rotated relative to the cylinder and the sample, like in Figure 6.2. These modifications do not incur computational costs, since the same cylinder discretization and LTM can be reused.

Another interesting direction for future research is to replace the Spectralon coating of the cylinder with a paper-like coating, to further reduce the cost of the setup. There is no inherent limitation for doing this, since the BRDF of the coating is taken into account in the LTM. However, the reflectance of paper is lower than that of Spectralon, and it is also less diffuse. This will decrease the equilibrium radiance that the camera captures, and might necessitate the use of a more sensitive camera to avoid excessively long exposure times.

The proposed design can also be used to capture BRDF. In that case, the sample will be illuminated from the same side as the camera. The illuminating beam and the camera could be at $\theta_i = 0^\circ$ to the measured sample, which will require a beamsplitter. For other incident directions the sample will be illuminated off center. The incident angles will, however, be limited, since the cylinder will block many directions, depending on the length of the cylinder. A workaround would be to also rotate the sample together with the light source.

Shapes different from a cylinder can also be used. The cylinder is convenient because it can easily and accurately be fabricated. Conical frustums are also interesting, because they can easily be fabricated, and allow for a different viewing geometry, which can reflect more direct radiance towards the camera, but the angles to the sample will become steeper. An in-depth discussion on the measurement geometry for radial imaging can be found in [55].

In Section 5.6 we introduced a single bounce version of the proposed design. It is based on a long flat stripe made out of diffusing material, that, by rotating, describes a cylinder. Alternatively, the sample could be rotated, while the stripe stays fixed. At each ϕ_o position, it reflects many θ_o angles. This implementation will reduce the computational requirements to a simple vector division: evaluating equation 5.58, where E_0 is the radiance captured by the camera divided by T_1 , or the BRDF for the light path sample \rightarrow pixel location \rightarrow camera. On the other hand, it would require many more captured sets of images. For example, we had 64 facets along ϕ_o in our experiments with the cylinder. The upper limit for the angular resolution along θ_o for this design is the resolution of the camera. The ideal camera for such a design would be an HDR push broom video camera.





Figure 6.3 – Schema of the mirrored cylinder. The lines of different color represent different number of bounces, blue is a single bounces, green is two bounces, and red is three bounces. The purple circles show the positions of the last light bounce. These are the positions that the camera registers in the image. The noted angles correspond to θ_o of the outgoing light directions.

Drawing another idea from the same work, [55], is to place mirrored rings inside the cylinder on the places where the camera observes the sample after one, two, three, or more bounces. These locations can easily be computed, and are shown on Figure 6.3 for the given geometry. Two images from two different measurement geometries are shown in Figure 6.4. Figure 6.4a shows the reflections of a sample with a radius of 2.5 mm, and the Figure 6.4b, shows the reflections of a sample with a radius of 1 mm, but also after the camera has moved closer to the cylinder and a wider angle lens has been used. The estimated BTDF at these positions is very accurate, and can be used as a guide for the BTDF of the other outgoing directions.



(a) Sample radius = 2.5 mm
Cylinder length = 2500 mm
Cylinder radius = 500 mm
Camera distance = 3000 mm
Focal length = 35 mm



(b) Sample radius 1 mm
Cylinder length = 2500 mm
Cylinder radius = 500 mm
Camera distance = 2600 mm
Focal length = 10 mm

Figure 6.4 – Simulated images for a diffuse sample and mirrored cylinders. The camera registers radiance only at the places where it see the sample after a certain number of bounces. The θ_o angle from the sample increases with the radius of the circle.

Bibliography

- [1] Andersen, M. (2004). Innovative bidirectional video-goniophotometer for advanced fenestration systems. Technical report, EPFL.
- [2] Andersen, M. and de Boer, J. (2006). Goniophotometry and assessment of bidirectional photometric properties of complex fenestration systems. *Energy and buildings*, 38(7):836–848.
- [3] Andersen, M., Stokes, E., Gayeski, N., and Browne, C. (2010). Using digital imaging to assess spectral solar-optical properties of complex fenestration materials: A new approach in video-goniophotometry. *Solar Energy*, 84(4):549–562.
- [4] Anderson, M., Motta, R., Chandrasekar, S., and Stokes, M. (1996). Proposal for a standard default color space for the internet - sRGB. In *IS&T Color and Imaging Conference (CIC'96)*, pages 238–245.
- [5] Apian-Bennewitz, P. (1994). Designing an apparatus for measuring bidirectional reflection/transmission. In *Optical Materials Technology for Energy Efficiency and Solar Energy Conversion XIII*, volume 2255, pages 697–706. International Society for Optics and Photonics.
- [6] Arecchi, A. V., Koshel, R. J., and Messadi, T. (2007). *Field guide to illumination*. SPIE Publications.
- [7] Ashdown, M. and Sato, Y. (2005). Steerable projector calibration. In *IEEE Computer Vision and Pattern Recognition Workshop (CVPRW'05)*, pages 98–98.
- [8] Audet, S. and Okutomi, M. (2009). A user-friendly method to geometrically calibrate projector-camera systems. In *IEEE Computer Vision and Pattern Recognition Workshop (CVPRW'09)*, pages 47–54.
- [9] Ben-Artzi, A., Egan, K., Durand, F., and Ramamoorthi, R. (2008). A precomputed polynomial representation for interactive brdf editing with global illumination. *ACM Transactions on Graphics (TOG)*, 27(2):1–13.
- [10] Ben-Ezra, M., Wang, J., Wilburn, B., Li, X., and Ma, L. (2008). An led-only brdf measurement device. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE.

Bibliography

- [11] Bhushan, B., Wyant, J. C., and Koliopoulos, C. L. (1985). Measurement of surface topography of magnetic tapes by mirau interferometry. *Applied Optics*, 24(10):1489–1497.
- [12] Blais, F., Picard, M., and Godin, G. (2004). Accurate 3d acquisition of freely moving objects. In *Proceedings. 2nd International Symposium on 3D Data Processing, Visualization and Transmission, 2004. 3DPVT 2004.*, pages 422–429. IEEE.
- [13] Bouguet, J.-Y. (2004). Camera calibration toolbox for matlab. http://www.vision.caltech.edu/bouguetj/calib_doc.
- [14] Brown, D. C. (1966). Decentering distortion of lenses. *Photogrammetric Engineering and Remote Sensing*, 32(3):444–462.
- [15] Chen, C.-Y. and Chien, H.-J. (2013). An incremental target-adapted strategy for active geometric calibration of projector-camera systems. *Sensors*, 13(2):2664–2681.
- [16] Chen, T., Seidel, H.-P., and Lensch, H. P. (2008). Modulated phase-shifting for 3d scanning. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE.
- [17] Dana, K. J. and Wang, J. (2004). Device for convenient measurement of spatially varying bidirectional reflectance. *JOSA A*, 21(1):1–12.
- [18] Debevec, P. E. and Malik, J. (1997). Recovering high dynamic range radiance maps from photographs. In *ACM International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'97)*, pages 369–378.
- [19] Deeb, R., Muselet, D., Hebert, M., and Tremeau, A. (2018). Interreflections in computer vision: a survey and an introduction to spectral infinite-bounce model. *Journal of Mathematical Imaging and Vision*, 60(5):661–680.
- [20] Deniel, J.-M. (2002). *Modélisation des luminaires et des BRDF: réalisation, mesure et compression*. PhD thesis, Rennes 1.
- [21] Dunn, W. (2011). Advertising displays. US Patent 8,016,452.
- [22] Dupuy, J. and Jakob, W. (2018). An adaptive parameterization for efficient material acquisition and rendering. *ACM Transactions on graphics (TOG)*, 37(6):1–14.
- [23] Elkhuisen, W. S., Lenseigne, B. A., Baar, T., Verhofstad, W., Tempelman, E., Geraedts, J. M., and Dik, J. (2015). Reproducing oil paint gloss in print for the purpose of creating reproductions of old masters. In *Measuring, Modeling, and Reproducing Material Appearance 2015*, volume 9398. International Society for Optics and Photonics.
- [24] Elkhuisen, W. S., Zaman, T., Verhofstad, W., Jonker, P. P., Dik, J., and Geraedts, J. M. (2014). Topographical scanning and reproduction of near-planar surfaces of paintings. In *Measuring, Modeling, and Reproducing Material Appearance*, volume 9018. International Society for Optics and Photonics.

-
- [25] Fleming, R. W. (2014). Visual perception of materials and their properties. *Vision research*, 94:62–75.
- [26] Geiger, A., Moosmann, F., Car, Ö., and Schuster, B. (2012). Automatic camera and range sensor calibration using a single shot. In *IEEE International Conference on Robotics and Automation (ICRA'12)*, pages 3936–3943.
- [27] Geometries, B. and Abrasion, S. (2012). Standard test method for haze and luminous transmittance of transparent plastics. *ASTM Int*, 1:1–7.
- [28] Ghosh, A., Achutha, S., Heidrich, W., and O'Toole, M. (2007). Brdf acquisition with basis illumination. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE.
- [29] Ghosh, A., Chen, T., Peers, P., Wilson, C. A., and Debevec, P. (2009). Estimating specular roughness and anisotropy from second order spherical gradient illumination. *Computer Graphics Forum*, 28(4):1161–1170.
- [30] Ghosh, A., Chen, T., Peers, P., Wilson, C. A., and Debevec, P. (2010). Circularly polarized spherical illumination reflectometry. *ACM SIGGRAPH Asia 2010 papers*, pages 1–12.
- [31] Goldman, D. B. (2010). Vignette and exposure calibration and compensation. *IEEE transactions on pattern analysis and machine intelligence*, 32(12):2276–2288.
- [32] Goodman, J. W. (2005). *Introduction to Fourier optics*. Roberts and Company Publishers.
- [33] gPhoto2 toolbox (2019). <http://www.gphoto.org>.
- [34] Grandin, K. and Roos, A. (1994). Evaluation of correction factors for transmittance measurements in single-beam integrating spheres. *Applied optics*, 33(25):6098–6104.
- [35] Griesser, A. and Van Gool, L. (2006). Automatic interactive calibration of multi-projector-camera systems. In *IEEE Computer Vision and Pattern Recognition Workshop (CVPRW'06)*, pages 8–8.
- [36] Guarnera, D., Guarnera, G. C., Ghosh, A., Denk, C., and Glencross, M. (2016). Brdf representation and acquisition. *Computer Graphics Forum*, 35(2):625–650.
- [37] Guarnera, G. C., Peers, P., Debevec, P., and Ghosh, A. (2012). Estimating surface normals from spherical stokes reflectance fields. In *European Conference on Computer Vision*, pages 340–349. Springer.
- [38] Gupta, M., Agrawal, A., Veeraraghavan, A., and Narasimhan, S. G. (2013). A practical approach to 3d scanning in the presence of interreflections, subsurface scattering and defocus. *International journal of computer vision*, 102(1-3):33–55.
- [39] Gupta, M. and Nayar, S. K. (2012). Micro phase shifting. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 813–820. IEEE.

Bibliography

- [40] Hahlweg, C. and Rothe, H. (2005). Design of a full-hemispherical spectro-radiometer with high dynamic range for characterization of surface properties using multispectral brdf data from vis to nir. In *Optical Fabrication, Testing, and Metrology II*, volume 5965. International Society for Optics and Photonics.
- [41] Hanssen, L. (2001). Integrating-sphere system and method for absolute measurement of transmittance, reflectance, and absorptance of specular samples. *Applied Optics*, 40(19):3196–3204.
- [42] Harris, C. and Stephens, M. (1988). A combined corner and edge detector. In *Citeseer Alvey Vision Conference (AVC'88)*, volume 15, page 50.
- [43] Hébert, M. and Hersch, R. D. (2011). Yule–nielsen based recto–verso color halftone transmittance prediction model. *Applied optics*, 50(4):519–525.
- [44] Hébert, M. and Machizaud, J. (2012). Spectral reflectance and transmittance of stacks of nonscattering films printed with halftone colors. *JOSA A*, 29(11):2498–2508.
- [45] Hsia, J. J. and Richmond, J. C. (1976). Bidirectional reflectometry. part i.: A high resolution laser bidirectional reflectometer with results on several optical coatings. *Journal of Research of the National Bureau of Standards. Section A, Physics and Chemistry*, 80(2):189.
- [46] Hurtos, T., Falcao, G., and Massich, J. (2008). Plane-based calibration of a projector camera system. *VIBOT Master*, pages 1–12.
- [47] ISO14782:1997 (1997). Plastics - determination of haze of transparent materials. Standard, International Organization for Standardization, Geneva, CH.
- [48] ISO25178-2:2012(en) (2012). Geometrical product specifications (gps) — surface texture: Areal — part 2: Terms, definitions and surface texture parameters. Standard, International Organization for Standardization, Geneva, CH.
- [49] Jakob, W. (2010). Mitsuba renderer. <http://www.mitsuba-renderer.org>.
- [50] JISK7361-1:1997 (1997). Plastics - determination of the total luminous transmittance of transparent materials - part 1: Single beam instrument. Standard, Japanese Industrial Standard, Tokyo, JP.
- [51] Karamata, B. and Andersen, M. (2013a). A review and analysis of parallel goniophotometry. *LUXEUROPA 2013*, 12(CONF).
- [52] Karamata, B. and Andersen, M. (2013b). Revisiting parallel catadioptric goniophotometers. In *Optical Measurement Systems for Industrial Inspection VIII*, volume 8788. International Society for Optics and Photonics.
- [53] Karamata, B. and Andersen, M. (2014). Origin and nature of measurement bias in catadioptric parallel goniophotometers. *JOSA A*, 31(5):1040–1048.

-
- [54] Kawamoto, H. (2002). The history of liquid-crystal displays. *Proceedings of the IEEE*, 90(4):460–500.
- [55] Kuthirummal, S. and Nayar, S. K. (2006). Multiview radial catadioptric imaging for scene capture. *ACM SIGGRAPH 2006 Papers*, pages 916–923.
- [56] Levoy, M., Pulli, K., Curless, B., Rusinkiewicz, S., Koller, D., Pereira, L., Ginzton, M., Anderson, S., Davis, J., Ginsberg, J., et al. (2000). The digital michelangelo project: 3d scanning of large statues. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 131–144.
- [57] Li, Z., Shi, Y., Wang, C., and Wang, Y. (2008). Accurate calibration method for a structured light system. *Optical Engineering*, 47(5).
- [58] Lin, H., Kim, S. J., Ssstrunk, S., and Brown, M. S. (2011). Revisiting radiometric calibration for color computer vision. In *IEEE International Conference on Computer Vision (ICCV'11)*, pages 129–136.
- [59] Machizaud, J. and Hbert, M. (2012). Spectral reflectance and transmittance prediction model for stacked transparency and paper both printed with halftone colors. *JOSA A*, 29(8):1537–1548.
- [60] Malzbender, T., Gelb, D., and Wolters, H. (2001). Polynomial texture maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 519–528.
- [61] Marschner, S. R., Westin, S. H., Lafortune, E. P., Torrance, K. E., and Greenberg, D. P. (1999). Image-based brdf measurement including human skin. In *Rendering Techniques' 99*, pages 131–144. Springer.
- [62] Mattison, P. R., Dombrowski, M. S., Lorenz, J. M., Davis, K. J., Mann, H. C., Johnson, P., and Foos, B. (1998). Handheld directional reflectometer: an angular imaging device to measure brdf and hdr in real time. In *Scattering and Surface Roughness II*, volume 3426, pages 240–251. International Society for Optics and Photonics.
- [63] Matusik, W., Pfister, H., Brand, M., and McMillan, L. (2003). A data-driven reflectance model. *ACM Transactions on Graphics*, 22(3):759–769.
- [64] Mazauric, S., Hbert, M., Simonot, L., and Fournel, T. (2014). Two-flux transfer matrix model for predicting the reflectance and transmittance of duplex halftone prints. *JOSA A*, 31(12):2775–2788.
- [65] McCluney, W. R. (2014). *Introduction to radiometry and photometry*. Artech House.
- [66] McNeil, J. R. and Wilson, S. R. (1993). Two-dimensional optical scatterometer apparatus and process. US Patent 5,241,369.

- [67] Mitsunaga, T. and Nayar, S. K. (1999). Radiometric self calibration. In *IEEE Computer Vision and Pattern Recognition (CVPR'99)*, volume 1.
- [68] Moreno, D. and Taubin, G. (2012). Simple, accurate, and robust projector-camera calibration. In *IEEE 3D Imaging, Modeling, Processing, Visualization & Transmission (3DIMPVT'12)*, pages 464–471.
- [69] Naik, N., Zhao, S., Velten, A., Raskar, R., and Bala, K. (2011). Single view reflectance capture using multiplexed scattering and time-of-flight imaging. In *Proceedings of the 2011 SIGGRAPH Asia Conference*, pages 1–10.
- [70] Nayar, S. K., Fang, X.-S., and Boulton, T. (1997). Separation of reflection components using color and polarization. *International Journal of Computer Vision*, 21(3):163–186.
- [71] Nayar, S. K., Ikeuchi, K., and Kanade, T. (1991). Shape from interreflections. *International Journal of Computer Vision*, 6(3):173–195.
- [72] Nayar, S. K., Peri, H., Grossberg, M. D., and Belhumeur, P. N. (2003). A projection system with radiometric compensation for screen imperfections. In *IEEE International Conference on Computer Vision Workshop (ICCV PROCAMS'03)*, volume 3.
- [73] Nicodemus, F. E. (1977). Self-study manual on optical radiation measurements: Part 1—concepts, chapters 1 to 3.
- [74] Nicodemus, F. E., Richmond, J. C., Hsia, J. J., Ginsberg, I. W., Limperis, T., et al. (1977). *Geometrical considerations and nomenclature for reflectance*, volume 160. Citeseer.
- [75] Nimier-David, M., Speierer, S., Ruiz, B., and Jakob, W. (2020). Radiative backpropagation: An adjoint method for lightning-fast differentiable rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)*, 39(4).
- [76] O'Toole, M. (2016). *Optical linear algebra for computational light transport*. PhD thesis, University of Toronto (Canada).
- [77] Page, M. (2018). *Création d'objets mats: optimisation d'un procédé d'impression en relief en termes d'apparence*. PhD thesis, Paris, CNAM.
- [78] Phan Van Song, T., Andraud, C., and Ortiz-Segovia, M. V. (2017). Spectral predictions of rough ink layers using a four-flux model. *Color and Imaging Conference*, 2017(25):251–257.
- [79] Pharr, M., Jakob, W., and Humphreys, G. (2016). *Physically based rendering: From theory to implementation*. Morgan Kaufmann.
- [80] Posdamer, J. L. and Altschuler, M. (1982). Surface measurement by space-encoded projected beam systems. *Computer graphics and image processing*, 18(1):1–17.
- [81] Ren, J. and Zhao, J. (2010). Measurement of a bidirectional reflectance distribution and system achievement based on a hemi-parabolic mirror. *Optics letters*, 35(9):1458–1460.

-
- [82] Rosete-Aguilar, M., Herrera, O. R., and Bruce, N. C. (2003). Optical design of a scatterometer with an ellipsoidal mirror. *Optical Engineering*, 42(6):1772–1778.
- [83] Rump, M., Müller, G., Sarlette, R., Koch, D., and Klein, R. (2008). Photo-realistic rendering of metallic car paint from image-based measurements. *Computer Graphics Forum*, 27(2):527–536.
- [84] Saint-Pierre, D., Chavel, P., Simonot, L., and Hébert, M. (2019a). Angular reflectance model for ridged specular surfaces, with comprehensive calculation of inter-reflections and polarization. *JOSA A*, 36(11):C51–C61.
- [85] Saint-Pierre, D., Deeb, R., Muselet, D., Simonot, L., and Hébert, M. (2018). Light interreflections and shadowing effects in a lambertian v-cavity under diffuse illumination. *Electronic Imaging*, 2018(8):166–1.
- [86] Saint-Pierre, D., Simonot, L., and Hébert, M. (2019b). Reflectance computation for a specular only v-cavity. In *International Workshop on Computational Color Imaging*, pages 289–303. Springer.
- [87] Salvi, J., Fernandez, S., Pribanic, T., and Llado, X. (2010). A state of the art in structured light patterns for surface profilometry. *Pattern Recognition*, 43(8):2666–2680.
- [88] Seitz, S. M., Matsushita, Y., and Kutulakos, K. N. (2005). A theory of inverse light transport. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 1440–1447. IEEE.
- [89] Sen, P., Chen, B., Garg, G., Marschner, S. R., Horowitz, M., Levoy, M., and Lensch, H. P. (2005). Dual photography. *ACM SIGGRAPH 2005 Papers*, pages 745–755.
- [90] Shahpaski, M., Ricardo Sapaico, L., Chevassus, G., and Süsstrunk, S. (2017). Simultaneous geometric and radiometric calibration of a projector-camera pair. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4885–4893.
- [91] Shahpaski, M., Sapaico, L. R., and Süsstrunk, S. (2019). Comparative analysis of transmittance measurement geometries and apparatus. *Electronic Imaging*, 2019(6):477–1.
- [92] Stavenga, D., Leertouwer, H., Pirih, P., and Wehling, M. (2009). Imaging scatterometry of butterfly wing scales. *Optics Express*, 17(1):193–202.
- [93] Su, X. and Zhang, Q. (2010). Dynamic 3-d shape measurement method: a review. *Optics and Lasers in Engineering*, 48(2):191–204.
- [94] Symons, J. G., Christie, E. A., and Peck, M. (1982). Integrating sphere for solar transmittance measurement of planar and nonplanar samples. *Applied optics*, 21(15):2827–2832.
- [95] Travis, J., Winchester, N., and Smith, M. (1995). Determination of the transmittance uniformity of optical filter standard reference materials. *Journal of research of the National Institute of Standards and Technology*, 100(3):241.

- [96] Tunwattanapong, B., Fyffe, G., Graham, P., Busch, J., Yu, X., Ghosh, A., and Debevec, P. (2013). Acquiring reflectance and shape from continuous spherical harmonic illumination. *ACM Transactions on graphics (TOG)*, 32(4):1–12.
- [97] Veach, E. (1997). *Robust Monte Carlo methods for light transport simulation*, volume 1610. Stanford University PhD thesis.
- [98] Ward, G. J. (1992). Measuring and modeling anisotropic reflection. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 265–272.
- [99] Wilm, J., Olesen, O. V., and Larsen, R. (2014). Accurate and simple calibration of DLP projector systems. In *SPIE MOEMS-MEMS*.
- [100] Wold, J. P., Jakobsen, T., and Krane, L. (1996). Atlantic salmon average fat content estimated by near-infrared transmittance spectroscopy. *Journal of Food Science*, 61(1):74–77.
- [101] Yamazaki, S., Mochimaru, M., and Kanade, T. (2011). Simultaneous self-calibration of a projector and a camera using structured light. In *IEEE Computer Vision and Pattern Recognition Workshop (CVPRW'11)*, pages 60–67.
- [102] Yu, H., Hsiao, C., and Liu, W. (2006). New apparatus for haze measurement for transparent media. *Measurement Science and Technology*, 17(8).
- [103] Yu, H.-L. and Hsiao, C.-C. (2009). Comparison of different measurement methods for transmittance haze. *Metrologia*, 46(4).
- [104] Zaman, T. (2013). *Development of a topographical imaging device for the near-planar surfaces of paintings*. PhD thesis, Delft University of Technology.
- [105] Zhang, S. and Huang, P. S. (2006). Novel method for structured light system calibration. *Optical Engineering*, 45(8).
- [106] Zhang, Z. (2000). A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 22(11):1330–1334.

MARJAN SHAHPASKI

Lausanne, Switzerland

🏠 Homepage  LinkedIn

EDUCATION

EPFL (Swiss Federal Institute of Technology in Lausanne) *September 2014 - October 2020*
PhD in Computer Science, advised by Prof. Sabine Süsstrunk *Lausanne, Switzerland*
Thesis title: *"Precise and accurate image-based measurements of optical material properties"*

EPFL (Swiss Federal Institute of Technology in Lausanne) *September 2012 - August 2014*
Master's degree in Computer Science, GPA: 5.51/6.00 *Lausanne, Switzerland*
Thesis title: *"3D color halftoning and prediction"*

FEEIT (Faculty of Electrical Engineering and Information Technologies) *Sep. 2006 - Feb. 2011*
Bachelor's degree in Computer Engineering, GPA: 9.31/10.00 *Skopje, Macedonia*
Thesis title: *"Automated stitching of multiple images into a panoramic image"*

PROFESSIONAL EXPERIENCE

EPFL, Image and Visual Representation Lab *September 2014 - October 2020*
Ph.D Student / Research Assistant *Lausanne, Switzerland*

- Developed a novel method for projector-camera pair calibration for structured light 3D scanners [6].
- Published research on accurate total transmittance measurement [1].
- Proposed a closed-form solution to the depth ambiguity in single-perspective images [3, 4].
- Designed a novel method for measuring BTDF based on radial catadioptric imaging (under submission).
- Taught courses in Computer Vision, Computational Photography, Distributed Algorithms and Physics. Won a Teaching Assistant Award.
- Supervised and mentored Bachelor's and Master's students.
- Multiple well-received presentations at EPFL, our industrial partner and at academic conferences.
- EPFL EDIC Fellowship for the first year of PhD studies in 2014/2015.

Disney Research *June 2017 - August 2017*
Research Assistant *Zürich, Switzerland*

- Developed an approach that produces striking changes in the appearance of prints when the illumination conditions are modified in a predetermined fashion.
- Our work resulted in a publication [2], and a patent [8].

Océ PLT - A Canon Company *September 2013 - February 2014*
Research Assistant *Paris, France*

- Developed methods that optimize the print quality of special effects, such as lenticular printing and mosaic creation, for the "Océ 3D high resolution printing" systems.
- Due to their efficacy, the methods have been included in the print controllers.
- Part of our work resulted in a publication [7], and a patent [9].

EPFL, Peripheral Systems Lab *June 2012 - August 2013*
Research Assistant *Lausanne, Switzerland*

- Developed an application for accurate preview of images that will be printed with daylight fluorescent inks for gamut extension. The application relies on CGAL for high execution speeds, and is cross-platform compatible.

TECHNICAL STRENGTHS

General	Computer vision, computational photography, computer graphics, image processing, color science, machine learning, deep learning, big data, data analysis, data mining, data visualization, algorithms, data structures, distributed algorithms, distributed information systems, GPU computing
Languages / Tools	Python, C, C++, MATLAB, R, Java, PyTorch, Hadoop, SQL, bash, \LaTeX , Git, Mitsuba, Windows, Linux, MacOS, Office suite

LANGUAGES

- Macedonian — Native
- English — Full professional proficiency
- French — Basic proficiency

EXTRACURRICULAR ACTIVITIES

- Computer hardware
- Biking
- Finance
- Photography, camera crew for a short film (“Les bruits des pas”)

PERSONAL INFORMATION

- Civil status: Married
- Swiss driving license: type B
- Nationality: Macedonian
- Swiss permit: type B

PUBLICATIONS AND PATENTS

- [1] Marjan Shahpaski, Luis Ricardo Sapaico, and Sabine Süsstrunk. Comparative analysis of transmittance measurement geometries and apparatus. *Electronic Imaging*, 2019.
- [2] Petar Pjanic, Marjan Shahpaski, and Anselm Grundhöfer. Magic prints: Image-changing prints observed under visible and 365 nm uv light. *Journal of Imaging Science and Technology*, 63(2), 2019.
- [3] Majed El Helou, Marjan Shahpaski, and Sabine Süsstrunk. Solving the depth ambiguity in single-perspective images. *OSA Continuum*, 2(10), 2019.
- [4] Majed El Helou, Marjan Shahpaski, and Sabine Süsstrunk. Closed-form solution to disambiguate defocus blur in single-perspective images. In *Mathematics in Imaging*. Optical Society of America, 2019.
- [5] Boaz Arad, Ohad Ben-Shahar, and Radu Timofte. Ntire 2018 challenge on spectral reconstruction from rgb images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2018.
- [6] Marjan Shahpaski, Luis Ricardo Sapaico, Gaspard Chevassus, and Sabine Süsstrunk. Simultaneous geometric and radiometric calibration of a projector-camera pair. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [7] Teun Baar, Marjan Shahpaski, and Maria V. Ortiz Segovia. Image ghosting reduction in lenticular relief prints. In *Measuring, Modeling, and Reproducing Material Appearance*. International Society for Optics and Photonics, 2014.
- [8] Petar Pjanic, Anselm Grundhöfer, and Marjan Shahpaski. Generating prints with multiple appearances. US Patent 10,491,784.
- [9] Teun Baar, Marjan Shahpaski, and Maria V. Ortiz Segovia. Ghosting compensation in relief images for directional prints. US Patent 9,418,323.