



AVANÇOS NA REPRESENTAÇÃO DE MATERIAIS PARA PROTOTIPAGEM VIRTUAL PREDITIVA

Olaf Daniel Clausen

Tese de Doutorado apresentada ao Programa de Pós-graduação em Engenharia de Sistemas e Computação, COPPE, da Universidade Federal do Rio de Janeiro, como parte dos requisitos necessários à obtenção do título de Doutor em Engenharia de Sistemas e Computação.

Orientador: Claudio Esperança

Rio de Janeiro
Novembro de 2022

AVANÇOS NA REPRESENTAÇÃO DE MATERIAIS PARA PROTOTIPAGEM
VIRTUAL PREDITIVA

Olaf Daniel Clausen

TESE SUBMETIDA AO CORPO DOCENTE DO INSTITUTO ALBERTO LUIZ
COIMBRA DE PÓS-GRADUAÇÃO E PESQUISA DE ENGENHARIA (COPPE)
DA UNIVERSIDADE FEDERAL DO RIO DE JANEIRO COMO PARTE DOS
REQUISITOS NECESSÁRIOS PARA A OBTENÇÃO DO GRAU DE DOUTOR EM
CIÊNCIAS EM ENGENHARIA DE SISTEMAS E COMPUTAÇÃO.

Examinada por:

Prof. Claudio Esperança, Ph.D.

Prof. Arnulph Fuhrmann, Dr.

Prof. Ricardo Guerra Marroquim, D.Sc.

Prof. Michael Weinmann, Dr.

Prof. Alexander Wilkie, Ph.D.

Prof. Manuel Menezes de Oliveira Neto, Ph.D.

Prof. Ricardo Farias, Ph.D.

RIO DE JANEIRO, RJ – BRASIL
NOVEMBRO DE 2022

Clausen, Olaf Daniel

Avanços na Representação de Materiais para Prototipagem Virtual Preditiva/Olaf Daniel Clausen. – Rio de Janeiro: UFRJ/COPPE, 2022.

XIX, 103 p.: il.; 29,7cm.

Orientador: Claudio Esperança

Tese (doutorado) – UFRJ/COPPE/Programa de Engenharia de Sistemas e Computação, 2022.

Bibliography: p. 94 – 103.

1. Virtual Prototyping. 2. Realistic Rendering. 3. Reflectance Modelling. 4. Diffraction. 5. Display Calibration. I. Esperança, Claudio. II. Universidade Federal do Rio de Janeiro, COPPE, Programa de Engenharia de Sistemas e Computação. III. Título.

Resumo da Tese apresentada à COPPE/UFRJ como parte dos requisitos necessários para a obtenção do grau de Doutor em Ciências (D.Sc.)

AVANÇOS NA REPRESENTAÇÃO DE MATERIAIS PARA PROTOTIPAGEM VIRTUAL PREDITIVA

Olaf Daniel Clausen

Novembro/2022

Orientador: Claudio Esperança

Programa: Engenharia de Sistemas e Computação

Prototipagem virtual (PV) é uma parte estabelecida e cada vez mais importante do processo de desenvolvimento de novos produtos. No entanto, PV ainda não é utilizado para certificar as etapas de desenvolvimento virtual. Esta limitação impede o desenvolvimento holístico e paralelo de produtos, onde engenheiros e designers dos parceiros mundialmente envolvidos desenvolvem simultaneamente um novo produto. Conseqüentemente, também impede uma redução considerável do tempo e dos custos de desenvolvimento, do uso de recursos e do impacto ambiental. A principal questão da PV é o limite da fidelidade dos protótipos virtuais, o que é particularmente problemático quando a PV é usada para decisões de design. Neste contexto, o protótipo virtual deve evocar a mesma sensação visual do que o protótipo físico; portanto, deve ser preditivo. Todavia, devido à complexidade da simulação de luz e às capacidades da visão humana, os protótipos virtuais muitas vezes não são preditivos. Dois grandes problemas são a representação do material e a reprodução de cores em um dispositivo de visualização. Nesta tese, focamos nestes dois problemas abordando os seguintes três tópicos:

Primeiro, apresentamos um novo framework de validação, que permite a validação da simulação da luz comparando-a com a realidade. A implementação com o renderizador científico Mitsuba confirmou que a modelagem da interação luz-matéria é a principal fonte de erro. Além disso, detectamos mudanças de comprimento de onda no comportamento de reflexos de amostras de cor, um fenômeno que é novo para a comunidade de Computação Gráfica.

Em segundo lugar, examinamos mais especificamente esses efeitos de mudança de comprimento de onda que levam a aparências avermelhadas e azuladas. Estes efeitos baseados em difração têm sido pouco relatados na literatura. Esta tese fornece a primeira análise completa a partir de dados coletados com precisão e apresenta uma extensão simples ao modelo Cook-Torrance considerando tais mudanças de comprimento de onda.

Finalmente, lidamos com a consistência de cores dos dispositivos de realidade virtual (Virtual Reality Head Mounted Displays - VR HMDs). As cores precisam ser apresentadas de forma consistente para aumentar a aceitação destes dispositivos na PV. Nesta tese, apresentamos uma abordagem para caracterizar e calibrar os monitores de VR HMDs de consumo. Combinado com nossa abordagem de calibração, é possível reproduzir cores com o dispositivo Vive Pro sem diferenças perceptíveis.

Abstract of Thesis presented to COPPE/UFRJ as a partial fulfillment of the requirements for the degree of Doctor of Science (D.Sc.)

TOWARDS PREDICTIVE VIRTUAL PROTOTYPING: IMPROVING MATERIAL REPRESENTATION FOR REALISTIC RENDERING

Olaf Daniel Clausen

November/2022

Advisor: Claudio Esperança

Department: Systems Engineering and Computer Science

Virtual prototyping (VP) is an established and increasingly important part of the development process of new products. However, VP is still not used to release virtual development stages. This limitation prevents holistic and parallel product development, where engineers and designers of all involved partners worldwide simultaneously develop a new product. Consequently, it hinders a considerable reduction of development time and costs, resource usage, and environmental impact. The main issue of VP is the limited fidelity of virtual prototypes, which is particularly problematic when VP is used for design decisions. In this application, the virtual prototype has to evoke the same visual sensation as the physical prototype; thus has to be predictive. However, due to the light simulation's complexity and human vision's capacities, virtual prototypes often fail to be predictive. Two major problems are the material representation and the colour reproduction on a display device. In this thesis, we focus on these two problems by tackling the following three topics:

First, we introduce a new validation framework, which enables the validation of the light simulation by comparing it with reality. The implementation with the scientific renderer Mitsuba confirmed that modelling the light-matter interaction is the main source of error. Moreover, we detected wavelength shifts in the reflections behaviour of colour patches, a phenomenon that is new to the CG community.

Second, we look more specifically into these wavelength shift effects that lead to reddish and blueish appearances. These diffraction-based effects have been

scarcely reported in the literature. This thesis provides the first thorough analysis from precise measured data and introduces a simple extension to the Cook-Torrance model that considers such wavelength shifts.

Finally, we deal with the colour consistency of virtual reality head mounted displays (VR HMDs). The colour consistency needs to be improved to increase the acceptance of HMDs as VR-System for VP. In this thesis, we present an approach to characterize and calibrate displays of consumer VR HMDs. Combined with our calibration approach, the Vive Pro reproduces colours without perceivable differences.

Acknowledgements

First of all, I would like to thank my three supervisors Ricardo Marroquim (TU Delft), Arnulph Fuhrmann (TH Köln) and Claudio Esperança (UFRJ). I was lucky enough to be supervised by three professors, who supported me in different ways; helping me with professional, scientific, personal, and bureaucratic problems. Your tireless support helped me to overcome the recurring downs during my doctoral project and to finally successfully complete my PhD.

I am very grateful to Holger Weigand for the countless discussions about optics topics and measurement results. I thank Klaus Dollinger († 2020), who primarily developed the gonioreflectometer, which was integral to my doctoral project. Further, he taught me to solve applied optics problems in unconventional ways. I appreciated the valuable advices and practical experiences of Aaron Finkenthei, which helped me to design and build components for the custom build measurement apparatus. I really enjoyed working together with Martin Miśiak and Kristoffer Waldow. You both created a supportive and non-competitive atmosphere in the office, which I always enjoyed. Thanks for helping me with programming and Blender issues and tolerating my unusual habits, e.g. doing a routine power nap on the sofa.

Finally, I thank my great family. First, I thank my mother and four siblings, who build a solid rock in my life, on which I can rely on whenever I need it. In particular, I thank my mother for helping us a lot with our daughter, my sister Adele for improving my scientific writing and my brother Hans for the endless research discussions and your interest in my work. Second, I thank my wonderful wife Martha and my two daughters Shaya and Anya. I really appreciate that you supported me during the last five years and endured my ups and downs with a lot of patience and without complaining. You three taught me to tackle challenges in a more relaxed and confident way. Te amo!

To my parents Cordula and Harald (+ 2012)

Contents

Acknowledgements	viii
List of Tables	xii
List of Figures	xiii
1 Introduction	1
1.1 Problem	2
1.2 Contributions	4
1.3 Organization of the thesis	5
2 Background	6
2.1 Optics	6
2.1.1 Radiometry	6
2.1.2 Colourimetry	8
2.1.3 Fresnel equations	13
2.1.4 Reflection functions	14
2.2 Gonioreflectometer	21
2.2.1 Setup	21
2.2.2 Calibration	23
2.2.3 Cross-validation	25
2.2.4 Detector validation	28
2.3 Other measurement devices	28
3 Validation of light simulation	31
3.1 Related work	32
3.2 Validation framework	34
3.3 Implementation	35
3.3.1 Normbox	35
3.3.2 Comparison path	36
3.3.3 Error path	40
3.4 Fidelity of ground truth data	42

3.5	Conclusion	51
4	Light interaction with rough surfaces	54
4.1	Related work	55
4.2	Experimental procedures	57
4.3	Observations	60
4.4	Analysis	63
4.5	Application	73
4.6	Conclusion	79
5	Colour in virtual reality	81
5.1	Related work	82
5.2	Definition and verification of the display model	83
5.3	Calibration of HTC Vive Pro and Pimax 5k+	86
5.4	Conclusion	90
6	Conclusion	91
	Bibliography	94

List of Tables

4.1	Sandblasting conditions to generate the eight aluminium samples with varying roughness	58
4.2	Averaged dE 2000 between the measured in-plane BRDFs of the aluminium samples (top table) and ColorChecker patches (bottom table) and their respective fittings with the Cook-Torrance and our model. For all samples, our model outperforms the Cook-Torrance model.	77
5.1	Results of our display characterization	89

List of Figures

2.1	A taxonomy of selected reflectance models (According to [1]) . . .	15
2.2	The BRDF is defined as the reflected spectral radiance dL_o within the cone $d\Omega_o$ towards ω_o divided by the incident spectral irradiance E_i within the cone $d\Omega_i$ coming from ω_i . (According to [2]) . .	16
2.3	Classification of BRDF representations in a continuum diagram. .	17
2.4	Our custom-built gonioreflectometer (a) consists of a light system (b) producing a uniform and collimated illumination, a detector (c) combining a focusing detector head with a high-precision spectrometer, and a two-axis gonio-stage rotating the sample and detector head.	22
2.5	Spectralon BRDF at 632 nm and the incident angles 10° , 30° and 45° acquired by Levesque et al., Durell et al., Jakob and us.	26
2.6	Spectral BRDF of Spectralon at an incident angle of 80° and the reflection angles -50° , -60° , -70° and -80° acquired by Levesque et al., us and Jakob.	27
2.7	In the first detector setup (a), the sample is illuminated by collimated light. The scattered light is captured and coupled into a fibre cable by an achromatic lens. The fibre cable is connected to the spectrometer QE Pro. In the second detector setup (c), the achromatic lens is replaced by an integrating sphere from Thorlabs and the sample is illuminated by converged light. In both detector setups, the BRDF spectra include the linear wavelength shift at grazing angles (b,d).	29
3.1	Overview of our validation framework.	34
3.2	The Normbox consists of a wooden box and the spectroradiometer CS-1000A by Konica Minolta. An integrating sphere combined with an opal panel is used as a light source.	36
3.3	Error sources that occurs during the implementation of the comparison path of the validation framework	41

3.4	Comparison of the reference and predicted spectral radiance of Patches 2 (a) and 23 (b) and the respective NSE.	43
3.5	Investigation of the correlation between the sample radiance and the RMSE. A clear positive linear relationship with a correlation coefficient of 0.84 and a moderate linear relationship with a correlation coefficient of 0.62 can be observed for small and large sample angles, respectively.	45
3.6	Investigation of the correlation between the sample radiance and the NRMSE. There is no correlation for small sample angles; contrary, an exponential increase of NRMSE for decreasing radiance can be observed for large sample angles.	46
3.7	Error model: The absolute error is simulated by multiplying the radiance with the light composition and the two errors, err_d (blue) and err_i (red). err_d is caused by the fitting error of the patch; err_i by the fitting error of the grey wall.	47
3.8	Top row: NSE_{pred} for patches 19, 23 and 24 at sample angle 30° (left), and for patch 21 at sample angle 90° (right). Bottom row: respective NSE_{fit}	48
3.9	Measured and fitted BRDF of the gray wall for three different reflection angles. From left to right, the plots show BRDF spectra which are dominated by back-scattering, diffuse reflection, and specular reflection.	49
3.10	Color difference (DeltaE 2000) between reference and predicted data of all patches and sample angles.	50
3.11	Comparison to related work.	52
4.1	Top and 3D views of the samples' surface topographies measured with a confocal microscope. Except for Sample 6, the surfaces consist of randomly distributed holes caused by the glass beads of the sandblaster.	59
4.2	Reflectance distribution of Samples 1, 2, 4, and 6 at wavelength 600 nm. Sample 6, with a mirror finish, has a narrow specular peak. Instead, the rougher samples scatter more broadly and present an off-specular peak for large incident angles.	60

4.3	Eight aluminium samples with surfaces varying from smooth (top-left) to rough (bottom-right). For each sample, we show macro photos of forward and back-scattering on the left and a detailed view, as well as the average colour on the right. For all samples, colourful diffraction phenomena can be observed, which is particularly strong for rough surfaces where the appearance strongly deviates from the expected greyish of aluminium.	61
4.4	In-plane BRDF of Sample 1 (left) depicted in an RGB image and simulated results using the wave-optics model from Yan et al. [3]. x and y-axis are, respectively, reflected (θ_o) and incident (θ_i) angles. The reddish tone in the mirror reflection (centre of the forward scattering) changes continuously into a greyish-bluish tone in the back-scattering region. For the simulated data, the back-scattering effect is not contemplated.	62
4.5	On the left side, the measured spectral BRDFs of Sample 1 for forward and back-scattering at $\theta_i = 40^\circ$. On the right, the same for simulated scattering using the wave-optics approach from Yan et al. [3]. The measured BRDF spectra are linearly shifted towards long and short wavelengths in the forward and back-scattering. Even though the increasing trend is noticeable in the simulations, back-scattering is not contemplated.	63
4.6	First order of light diffraction at gratings with varying spacing (a) and the resulting spectra when binning over the diffraction angles (b). Diffraction is a reasonable explanation for the linear wavelength shifts observed on rough surfaces.	66
4.7	Comparison of the local and global fit with the measured data for Sample 1 using the RGB representation. On the left, the local fit is divided into the Fresnel and Diffraction term. The Delta E's 2000 between the fitted and the measured BRDFs are shown on the bottom row. In the case of local fit, the Delta E's remain below the perception threshold. The global fit leads to slightly higher Delta E's but can still adequately reproduce the measured data.	67
4.8	Slope distribution (y-axis) of Sample 1 resulting from the local fit (blue line) and the global fit using a cosine function (orange line).	68
4.9	Results of our global fitting model for Samples 2-5.	69
4.10	Simulation results show in blue the percentage of energy coming from single reflection rays and in light blue multi-reflection rays for wavelength 380 nm and varying incident angles.	70

4.11	Slope distributions locally fitted with a linear function of Spectralon SRS 02 sanded with grain size in the CAMI scale of 60, 1000, 2500. Back-scattering dominates for roughness 60, while forward scattering dominates for roughness 2500. For roughness 1000, forward and back-scattering compensate each other.	71
4.12	Comparison of the fitting of a conductor (a) and a dielectric (b) with the Cook-Torrance and our model.	76
4.13	Renderings computed with Mitsuba 2 [4] and the fitted BRDFs of aluminium samples (a) and ColorChecker patches (b). Clear differences between the Cook-Torrance and our model can be observed for both conductors and dielectrics, as confirmed by large Delta E 2000 values. Note the different scales of the colour bars.	78
5.1	The Display model.	83
5.2	Verification of the display model.	85
5.3	Verification of the display model application.	85
5.4	Display characterization of the HTC Vive Pro and Pimax 5k+. The left and right columns show the results for the left and right displays. (a) Black level. (b) Original white spectrum and the simulation with the display model. (c) Chromaticity diagram including the sRGB, AdobeRGB, and display colourspaces. (d) DeltaE histogram (see Figure 5.2).	87
5.5	Comparison of measured colours and its reproduction with the sRGB, the AdobeRGB, and the measured display profile on the Vive Pro.	89

List of Symbols

The next list describes several symbols that will be later used within the body of the document

(θ_i, ϕ_i)	Polar and azimuthal angle of incidence direction ω_i
(θ_o, ϕ_o)	Polar and azimuthal angle of reflection direction ω_o
(θ_t, ϕ_t)	Polar and azimuthal angle of refraction direction ω_o
$(L^* a^* b^*)$	CIE Uniform colourspace
(R, G, B)	nonlinear display control values
(r, g, b)	linear display control values
(x, y)	Chromaticity coordinates
(X, Y, Z)	CIE Tristimulus Values
$\bar{x}(\lambda), \bar{y}(\lambda), \bar{z}(\lambda)$	Colour matching functions
ΔE_{00}^*	CIE 2000 colour-difference Formula
ΔE_{ab}^*	CIE 1976 colour-difference Formula
η	Refractive index
λ	Wavelength
Ω	Solid angle
ω_h	Half-direction
ω_i	Direction from which light is incident
ω_o	Direction in which light is reflected
ω_t	Direction in which light is refracted
Φ	Radiant flux

C_λ	Calibration factor
$D(\omega_h)$	Normal distribution function
E	Irradiance
$F(\omega_i, \omega_h)$	Fresnel factor
$f_r(\lambda, \omega_i, \omega_o)$	Bidirectional reflectance distribution function
$G(\omega_o, \omega_i, \omega_h)$	Geometric attenuation factor
$I(\lambda)$	Spectral power distribution of standard illuminant
JND	Just noticeable difference
$k_{d,\lambda}$	Diffuse albedo
$k_{s,\lambda}$	Specular albedo
$L(\omega)$	Radiance along the ray in direction ω
$L(\omega, \lambda)$	Spectral radiance along the ray in direction ω
M	RGB- to XYZ-colourspace transformation matrix
n	Surface normal
$NRMSE$	Normalized root-mean-square error
NSE	Normalized spectral error
$P(\lambda)$	Spectral power distribution
Q	Radiant energy
R	Reflectance for unpolarized light
R_{\parallel}	Reflectance parallel to the plane of incidence
R_{\perp}	Reflectance perpendicular to the plane of incidence
$RMSE$	Root-mean-square error
$S(\lambda)$	Sample reflectance or transmission
V	Measured signal

1. Introduction

In the past, the development process of new products was straightforward and well-defined. It mainly focused on two factors, product function and costs. Companies had to offer a high-quality and -functional product for a fair price to be competitive. This development process led to products with a long life cycle and a small diversity.

In the last few decades, both factors have drastically changed. The product life shortened considerably while product diversity increased immensely. There are two main reasons for this development. First, apart from function and costs, customisation is an additional important factor during product development. And second, an increasing number of products depend on electronic components, which have to be updated in short cycles due to the rapid development of computer techniques. In the last few years, additional factors further complicated the development process. Due to climate change, customers are becoming more environmentally aware and demanding sustainable and environmentally friendly products. At the same time, the environmental standards are becoming stricter and the resources scarcer. These factors contradict the currently short product life cycle.

Satisfying these contradicting customers' needs is a big challenge. To meet this challenge, companies must shorten the product development process further, save resources and follow strict environmental standards while increasing product complexity and diversity. It is inevitable to redesign the development process so that a highly optimized product is available in a shorter time and at an earlier development stage than before while focusing more consistently than ever on customer needs.

Virtual engineering (VE) plays a key role in such a novel development process. VE describes single applications of virtual technologies in product development, contrary to virtual product development (VPD), a novel development process consistently based on virtual technologies. A widespread application of VE is virtual prototyping (VP). VP is used to optimize the product design, function and properties with physical simulations and system investigations based on computational-intensive data models [5]. In many companies, VP is already

integral to product development. It is particularly relevant in industrial areas with complex products, e.g. automotive, aerospace and architecture, but also in industrial branches providing products with strong dynamics of change, e.g. fashion.

Although VP is already being used successfully, it still has one major drawback that limits its benefits. Usually, virtual prototypes are not utilized as binding releases of development stages since they serve a purely informative purpose. Thus, virtual development stages still require physical prototypes as approval. As a result, most companies continue a sequential development process with several time-consuming and costly development cycles, where at the end of each cycle is a physical prototype. This old fashion development process prevents holistic and parallel product development, where engineers and designers of all partners worldwide simultaneously develop a product.

A novel development process based on virtual releases would considerably reduce the development time and costs, resource usage and environmental impact due to fewer required development cycles, physical prototypes and business meetings. Consequently, mastering the virtual process chains addresses all previously mentioned challenges in the product development process.

The question immediately arises: *Why are virtual prototypes not used as releases of virtual development stages?* The answer is simple: *The fidelity of virtual prototypes is insufficient to replace physical prototypes.*

1.1 Problem

As already mentioned, VP is used, among others, to optimize product design. In this application, the virtual prototype has to evoke the same visual sensation as the physical prototype. The appearance of a virtual prototype depends roughly on its geometric representation, the light simulation and the visualization on a display device.

Light simulation greatly impacts the appearance of a virtual prototype because it defines the overall colour of the virtual object. It is particularly challenging due to the complexity of the light-matter interaction and the energy transport of electromagnetic radiation. Both are well described in physics, but only for specific cases under defined boundary conditions. The scenes used in VP are much too complex to simulate light in a holistic physical way. Therefore, in computer graphics, the light simulation is heavily approximated and limited to mostly obeying a few fundamental laws of physics. These approximations simplify the light simulation by taking advantage of the limitations of human

vision so that it is computed within an appropriate time or even in real-time, ideally without introducing perceptible errors. However, in practice, the latter is often not respected, which can be traced back to two main problems. First, some physical phenomena are known but too complex to simulate in a reasonable time. Second, some perceptible phenomena are not well understood or even unknown. A good example applying to both problems is the in computer graphics widespread assumption that wave-optics phenomena are negligible. On the one hand, the assumption heavily simplifies the light simulation enabling real-time rendering. On the other hand, wave-optics phenomena, particularly on rough surfaces, are not well investigated; hence their influence on the material appearance is unknown. The deviations of rendered images from reality are an huge issue for color-critical applications such as VP, because they can lead to wrong design decisions. The limited and often unknown fidelity of virtual prototypes considerable limits their benefits.

A common way to determine and improve the fidelity of the light simulation is the comparison with measured data representing reality. Based on the comparison result, the light simulation's bottlenecks can be identified, further investigated by experiments and improved with new algorithms.

Another critical bottleneck of VP is the visualization of the virtual content on a display device. Although human vision has limitations that are used to simplify the light simulation, it is still a powerful sensory organ. Human eyes provide stereoscopic vision, high spatial- and temporal resolution, a large FOV and the ability to adapt to brightness and distance. In VP, different visualization systems are used, but none addresses all capacities of human eyes.

A popular visualization system for VP are virtual reality head mounted displays (VR HMDs). These devices track the user's head movements, provide stereoscopy and enable the user to interact with virtual objects. VR-HMDs have become more attractive to the industry with the advent of consumer devices, such as the Oculus Rift or the HTC Vive. These devices provide a reasonable display resolution, a large field of view, and accurate head tracking for a low price. It is to be expected that the high interest of big tech companies and science will considerably speed up the hardware and software development of HMDs, so their benefits progressively outweigh their limitations. There has already been extensive research on tracking, display resolution, and immersion for VR HMDs. However, little research has been done on display characterization and calibration from a colour perspective.

1.2 Contributions

This thesis has five major contributions, as listed below. The first four contributions are related to the light simulation and the last to the colour consistency of VR-HMDs.

Validation framework based on ground truth data (GTD)

We provide a validation framework based on acquired Ground Truth Data that enables the spectral validation of the light simulation by comparing it with reality. The GTD accurately describes a novel reference scene, including isotropic in-plane BRDFs of 24 colour patches and reference measurements of all patches under 13 angles captured inside the reference scene. Our reference data covers rough materials with different spectral distributions and under several illumination conditions, varying from direct to indirect light dominated conditions. The GTD has been published and is freely accessible under the following link:

<https://cg.web.th-koeln.de/brdf-gtd/>.

Validation of Mitsuba with GTD

We implement the validation framework with the scientific renderer Mitsuba. The spectral comparison of the simulation and reference data confirms the high fidelity of the GTD and the reflection model as the dominant source of error. Moreover, we detected wavelength shifts in the reflections behaviour of the colour patches, a phenomenon that is new to the computer graphics community.

Experimental investigation of diffraction on rough surfaces

We look more specifically at these wavelength shifts that lead to reddish and blueish appearances. These wavelength shifts have been scarcely reported in the literature, and in this thesis, we provide the first thorough analysis from precise measured data. A multi-modal dataset of eight aluminium samples with varying roughness has been published and is freely accessible under the following link:

<http://cg.web.th-koeln.de/mmd/>.

Diffraction extension to microfacet model

We propose an empirical model that can closely fit such phenomena, where the slope of the linear wavelength shift depends on the incident and reflection angles. Based on this model, we introduce an extension to the Cook-Torrance model that considers such wavelength shifts without affecting the computation time.

Characterisation and calibration of consumer VR HMDs

We present an approach to characterize and calibrate displays of consumer VR HMDs. The approach is based on a simple display model, commonly used for calibrating conventional displays, but has not yet been applied for VR HMDs. We implement this approach with the HTC Vive Pro and the Pimax 5k+. Combined with our calibration approach, the Vive Pro provides consistent colour reproduction without perceivable differences.

1.3 Organization of the thesis

The thesis is divided into a total of six chapters. It is organized as follows: In Chapter 2 we introduce basic radiometric and colourimetric quantities, the reflection function, more specifically, the bidirectional reflectance distribution function (BRDF), and all used measurement devices. The following three chapters are closely related to the papers [6], [7] and [8], which were submitted during the doctoral project. These papers are related but still tackle different problems; hence we give in each chapter a short introduction, related work and a conclusion. In Chapter 3 we present the novel validation framework for light simulation. The framework is based on Ground Truth Data, acquired in a novel reference scene called Normbox. The Normbox and the acquired GTD are explained in detail. Further, we implement the validation framework with the scientific renderer Mitsuba to determine the GTD fidelity by physically and perceptually comparing the predicted with the reference data. In Chapter 4 we experimentally investigate diffraction phenomena on rough surfaces by acquiring a multi-modal dataset of aluminium samples with varying roughness. The data are thoroughly analyzed and the observed diffraction phenomena are described in a heuristic model. This model proposes an extension to the popular Cook-Torrance GGX BRDF model. The model is verified by fitting it against the in-plane BRDF of the aluminium samples and colour patches. In Chapter 5, we present an approach to characterize and calibrate displays of consumer VR HMDs. We introduce a display model and two frameworks to virtually and experimentally verify the display model. We implement both with the consumer VR HMDs: HTC Vive Pro and Pimax 5k+. Finally, in Chapter 6, we summarise the results of all chapters and present research topics that emerge from this thesis.

2. Background

This thesis deals with the validation of light simulation, the investigation of light interaction with rough surfaces, and the calibration of VR HMDs. Since all three topics involve physical experiments primarily concerning with optics and colour science, we introduce in this chapter fundamental radiometric and colourimetric quantities, the Fresnel equations, and the reflection function, more precisely, the bidirectional reflectance distribution function (BRDF). Furthermore, we present the measurement instruments used and their measurement errors.

2.1 Optics

Optics is the branch of physics studying light propagation and the light interaction with matter, as well as the construction of instruments that use or detect it [9]. Light comprises electromagnetic radiation in the perceptible wavelength range of the human eyes, which ranges from 380 to 780 nm. Optic's laws are also valid for a broader electromagnetic spectrum, including ultraviolet and infrared light.

There are two classical approaches used in optics: geometric optics and wave-optics. Geometric optics assumes that the energy is transported along light rays, thus neglecting the wave characteristics of light. This approximation is valid as long as the component apertures and the surface roughness are larger than the light's wavelengths. Within these limits, the optical laws may be formulated with the language of geometry [10, p.116]. Otherwise, physical phenomena such as interference and diffraction occur, which can be explained only by the wave nature of light, thus by wave-optics.

2.1.1 Radiometry

Radiometry, in turn, is the branch of optics dealing with light measurement. The radiometric quantities, radiance and irradiance are fundamental in computer graphics and colourimetry. In the following, we briefly review the derivations of both quantities.

Solid angle

The solid angle in space is equivalent to the angle in planimetry. It is essential in radiometry due to the spherical propagation of radiation in space. The solid angle describes the ratio of a sphere's segment area A , usually the intersection of a cone with the sphere, to the square of the sphere's radius r . The quantity has the dimensionless unit steradian [sr].

$$\Omega = \frac{A}{r^2} \quad (2.1)$$

Radiant energy

Radiometry primarily deals with measuring the energy of electromagnetic radiation (radiant energy) and derived quantities. The SI unit of radiant energy is joule [J] and the symbol Q represents it.

Radiant flux

The radiant flux is the radiant energy per unit time emitted, reflected, transmitted, or received by an object. Radiant flux is specified in watts [$W = J \cdot s^{-1}$].

$$\Phi = \frac{dQ}{dt} \quad (2.2)$$

Irradiance

The quantity of irradiance is defined as the radiant flux per unit area. It usually measures the received radiant energy coming from the upward hemisphere. The definition is as follows:

$$E = \frac{d\Phi}{dA} \quad (2.3)$$

Irradiance has the units [$W \cdot m^{-2}$].

Radiance

Radiometers are widespread measurement instruments in radiometry and usually consist of a focusing optic and a detector. It integrates the radiant energy of a defined area within a solid angle and over a specific time, thus measuring radiance. Radiance is defined as:

$$L(\omega) = \frac{d^2\Phi(\omega)}{d\Omega(\omega) dA_{\omega}^{\perp}}, \quad (2.4)$$

where dA_{ω}^{\perp} is the differential measurement area dA projected towards ω and $d\Omega(\omega)$ is the differential solid angle around ω . The radiant flux $\Phi(\omega)$ counts the number of photons per unit time passing through the projected area within the solid angle. Radiance has the units $[W \cdot m^{-2} \cdot sr^{-1}]$.

Spectral radiance

Continuing with the previous example, radiometers extended by a diffraction grating measure the radiance for narrowed wavelength intervals. These devices are called spectroradiometers and measure the spectral radiance, which is defined as:

$$L(\omega, \lambda) = \frac{d^2\Phi(\omega)}{d\Omega(\omega) dA_{\omega}^{\perp} d\lambda}. \quad (2.5)$$

The quantity of spectral radiance has the units $[W \cdot m^{-2} \cdot sr^{-1} \cdot nm^{-1}]$.

2.1.2 Colourimetry

Strictly speaking, colourimetry is part of colour science and not optics. Nevertheless, we will review it at this point due to its close relation to optics.

Colourimetry describes the colour of a visual stimulus numerically. The basis of colourimetry is the trichromatic generalization that sums up the experimental laws of colour matching. A key statement of the trichromatic generalization is that “under a wide range of conditions of observations, many colour stimuli can be matched in colour completely by additive mixture of three fixed primary stimuli whose power have been suitable adjusted.”[11, p. 117] This statement implies that a three-dimensional vector (tristimulus values) together with three carefully chosen basic functions (colour-matching functions) describe colour stimuli. The required colour-matching functions are determined experimentally.

CIE tristimulus values

The International Commission on Illumination (CIE) introduced in 1931 and 1964 colour-matching functions for a 2° and 10° standard colourimetric observer. The 2° and 10° standard observers are recommended for viewing of matching fields of angular subtense from 1° to 4° and greater than 4° , respectively. These two sets of colour-matching functions are the basis of the applied colourimetry.

The tristimulus values XYZ of an emissive source with the spectral power

distribution $P(\lambda)$ are calculated as:

$$\begin{aligned} X &= \int_{\lambda} \bar{x}(\lambda) P(\lambda) d\lambda, \\ Y &= \int_{\lambda} \bar{y}(\lambda) P(\lambda) d\lambda, \\ Z &= \int_{\lambda} \bar{z}(\lambda) P(\lambda) d\lambda, \end{aligned} \quad (2.6)$$

where \bar{x} , \bar{y} , \bar{z} are the colour matching functions of the respective standard observer. The integration is taken over the light spectrum of 390 to 830 nm for the 2° and 360 to 830 nm for the 10° standard observer.

In practice, the colour-matching functions and the spectral power distribution are represented in discrete values. Thus, the integrals are replaced by summations.

$$\begin{aligned} X &= \sum_i \bar{x}_i(\lambda) P_i(\lambda) \Delta\lambda, \\ Y &= \sum_i \bar{y}_i(\lambda) P_i(\lambda) \Delta\lambda, \\ Z &= \sum_i \bar{z}_i(\lambda) P_i(\lambda) \Delta\lambda. \end{aligned} \quad (2.7)$$

In the case of reflective or transmissive samples, the spectral power distribution, $P(\lambda)$, is replaced by the sample reflectance or transparency, $S(\lambda)$, multiplied by the spectral power distribution of the source, $I(\lambda)$, illuminating the object.

$$\begin{aligned} X &= k \int_{\lambda} \bar{x}(\lambda) S(\lambda) I(\lambda) d\lambda, \\ Y &= k \int_{\lambda} \bar{y}(\lambda) S(\lambda) I(\lambda) d\lambda, \\ Z &= k \int_{\lambda} \bar{z}(\lambda) S(\lambda) I(\lambda) d\lambda. \end{aligned} \quad (2.8)$$

The factor k is a normalization factor that can be chosen arbitrarily, provided it is kept constant throughout any particular discussion. Usually, the factor is defined as $k = \frac{1}{\int_{\lambda} I(\lambda) \bar{y}(\lambda) d\lambda}$, hence the Y-tristimulus value of an ideal diffuse reflector is equal to one. Another common definition is $k_m = 683 \text{ lm/W}$, where the Y-tristimulus value becomes the luminance (photometric equivalent to radiance) in the CIE 1931 System, provided $I(\lambda)$ is measured in terms of spectral radiance. In practice, the integrals in Equation 2.8 are again replaced by summations as specified in the following:

$$\begin{aligned}
X &= k \sum_{\lambda} \bar{x}_i(\lambda) S(\lambda) I(\lambda) \Delta\lambda, \\
Y &= k \sum_{\lambda} \bar{y}_i(\lambda) S(\lambda) I(\lambda) \Delta\lambda, \\
Z &= k \sum_{\lambda} \bar{z}_i(\lambda) S(\lambda) I(\lambda) \Delta\lambda.
\end{aligned} \tag{2.9}$$

CIE xyY-colourspace

In colourimetry, it is more intuitively to describe a colour stimulus by its brightness and chromaticity instead of the tristimulus values. The chromaticity describes the colour's hue and is defined by the chromaticity coordinates, which are functions of the XYZ-tristimulus values:

$$\begin{aligned}
x &= \frac{X}{X + Y + Z}, \\
y &= \frac{Y}{X + Y + Z}.
\end{aligned} \tag{2.10}$$

The colour's brightness is described by the luminance, usually given by the Y-tristimulus value (see previous section).

RGB-colourspace

Conventional displays reproduce a given colour stimulus by overlaying three primary colours red, green and blue. Due to physical limitations, display devices always cover only a subspace of the XYZ-colourspace. This device-dependent subspace is called RGB-colourspace. The transformation from the RGB- to XYZ-colourspace is described by the 3×3 matrix M . The matrix is calculated with the chromaticity coordinates of the display primaries x_i and y_i with $i \in r, g, b$ and the tristimulus values of the display whitepoint $X_W Y_W Z_W$, usually a standard illuminant.

$$\begin{pmatrix} r \\ g \\ b \end{pmatrix} = M^{-1} \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} \tag{2.11}$$

$$M = \begin{pmatrix} S_r X_r & S_g X_g & S_b X_b \\ S_r Y_r & S_g Y_g & S_b Y_b \\ S_r Z_r & S_g Z_g & S_b Z_b \end{pmatrix} \tag{2.12}$$

$$\begin{aligned}
X_i &= \frac{x_i}{y_i} \\
Y_i &= 1 \\
Z_i &= \frac{1 - x_i - y_i}{y_i} \\
i &\in r, g, b
\end{aligned} \tag{2.13}$$

$$\begin{pmatrix} S_r \\ S_g \\ S_b \end{pmatrix} = \begin{pmatrix} X_r & X_g & X_b \\ Y_r & Y_g & Y_b \\ Z_r & Z_g & Z_b \end{pmatrix}^{-1} \begin{pmatrix} X_W \\ Y_W \\ Z_W \end{pmatrix} \tag{2.14}$$

The transfer function, also known as the tone response curve (TRC) or gamma, describes the non-linear reproduction of the display. The TRC is usually defined by a power function as follows:

$$\begin{aligned}
v &\in \{r, g, b\}, \\
V &\in \{R, G, B\},
\end{aligned} \tag{2.15}$$

$$V = v^{\frac{1}{\gamma}}. \tag{2.16}$$

Usually, displays are not calibrated because costly hard- and software, as well as experts, are required. Nevertheless, to enable sound colour reproduction, displays are designed to obey standard RGB-colourspaces, such as the widespread colourspaces sRGB and AdobeRGB. The following table provides the properties of both colourspaces.

colourspace	Chromaticity	Red	Green	Blue	Whitepoint
sRGB	x	0.6400	0.3000	0.1500	0.3127
	y	0.3300	0.6000	0.0600	0.3290
	Y	0.2126	0.7152	0.0722	1.0000
AdobeRGB	x	0.6400	0.3000	0.1500	0.3127
	y	0.3300	0.6000	0.0600	0.3290
	Y	0.2126	0.7152	0.0722	1.0000

The TRC of the AdobeRGB colourspace is computed as defined in Equation 2.16 with $\gamma = 2.2$. The sRGB colourspace uses a slightly different gamma

correction.

$$V = \begin{cases} 12.92 v & \text{if } v \leq 0.0031308 \\ 1.055 v^{\frac{1}{2.4}} & \text{otherwise} \end{cases}$$

$L^* a^* b^*$ -colourspace

A common problem in colourimetry is the numerical expression of the perceived colour difference between two given object-colour stimuli that mismatch in colour. Therefore, colour-difference formulas are of great importance. Referring to Wyszecki and Stiles [11, p. 165], none of the formulas proposed in the last few decades sufficiently solve this problem. However, the CIE recommends the perceptually almost uniform colourspaces $L^* u^* v^*$ and $L^* a^* b^*$. The $L^* a^* b^*$ -space with the associated colour-difference formula Delta E has been established in industry and science. The quantities L^*, a^*, b^* are defined by:

$$\begin{aligned} L^* &= 116 \left(\frac{Y}{Y_n} \right)^{1/3} - 16, \\ a^* &= 500 \left[\left(\frac{X}{X_n} \right)^{1/3} - \left(\frac{Y}{Y_n} \right)^{1/3} \right], \\ b^* &= 200 \left[\left(\frac{Y}{Y_n} \right)^{1/3} - \left(\frac{Z}{Z_n} \right)^{1/3} \right], \end{aligned} \quad (2.17)$$

where L^* describes the perceptual lightness, a^* the chromaticity on the green and red axis, and b^* the chromaticity on the blue and yellow axis. (X_n, Y_n, Z_n) are the tristimulus values of a diffuse white object illuminated by a reference light source. The L^*, a^*, b^* definition in Equation 2.17 is constrained to $X/X_n, Y/Y_n, Z/Z_n > 0.01$.

Colour-difference formulas

The perceived colour difference ΔE_{ab}^* between two colour stimuli is defined as the euclidean distance of these colours in the $L^* a^* b^*$ -space. The definition is as follows:

$$\Delta E_{ab}^* = [(\Delta L^*)^2 + (\Delta a^*)^2 + (\Delta b^*)^2]^{1/2}. \quad (2.18)$$

The ΔE_{ab}^* formula assumes that the $L^* a^* b^*$ -space is perceptually uniform, which is not the case. This non-uniformity leads to different interpretations of ΔE_{ab}^* values depending on location in the $L^* a^* b^*$ -space. Despite this issue, the ΔE_{ab}^* formula is still commonly used, probably due to its simplicity. Note that the just noticeable difference (JND) in a side-by-side comparison is 2.3 and not 1, as intended.

The CIE published a revised colour difference formula in 2000 called ΔE_{00}^* .

This formula is also based on the $L^* a^* b^*$ -space, but fixed some uniformity issues by implementing five corrections: adding a hue rotation term to deal with the problematic blue region and three terms to compensate for neutral colours, lightness, chroma, and hue. In this thesis, the weighting factors KL, KC, and KH are set to the default value. Since the derivation of the ΔE_{00}^* formula is long, we refer the reader to the CIE Tech Report [12] for the definition.

The ΔE_{00}^* values are classified by the rating scale given in the book Color Imaging [13, p.461]. The rating scale is actually defined for the ΔE_{94}^* -metric but can also be used for the ΔE_{00}^* due to their similarity. In the book, a Delta ΔE^* of one defines the JND in a side-by-side comparison, a value of 2 leads to discernible colour differences for patches next-to-each-other and values greater than five are easily perceived in a side-by-side comparison.

2.1.3 Fresnel equations

The Fresnel equations are fundamental for modelling the light-matter interaction in computer graphics. The Fresnel equations describe how a plane wave hitting an interface of two homogeneous media with different optical properties splits into a reflected wave that propagates back into the first medium and a refracted wave that proceeds into the second medium. The fraction of the incident plane wave reflected at the interface defines the quantity of reflectance. The reflectance depends on the incident wave's polarization and is defined for parallel and perpendicular polarized light as follows:

$$\begin{aligned} R_{\parallel} &= \left| \frac{\eta_2 \cos \theta_i - \eta_1 \cos \theta_t}{\eta_2 \cos \theta_i + \eta_1 \cos \theta_t} \right|^2, \\ R_{\perp} &= \left| \frac{\eta_1 \cos \theta_i - \eta_2 \cos \theta_t}{\eta_1 \cos \theta_i + \eta_2 \cos \theta_t} \right|^2, \end{aligned} \quad (2.19)$$

where η_1 and η_2 are the refractive indices of the first and second medium. θ_i and θ_t are the incident and transmission angles. Usually, in computer graphics, light rays are located in the air before hitting an object. Substituting the refractive index of air (≈ 1.0) for η_1 and Snell's law for the transmission angle θ_t results in:

$$\begin{aligned} R_{\parallel} &= \left| \frac{-\eta_2^2 \cos \theta_i + \sqrt{\eta_2^2 - \sin^2 \theta_i}}{\eta_2^2 \cos \theta_i + \sqrt{\eta_2^2 - \sin^2 \theta_i}} \right|^2, \\ R_{\perp} &= \left| \frac{\cos \theta_i - \sqrt{\eta_2^2 - \sin^2 \theta_i}}{\cos \theta_i + \sqrt{\eta_2^2 - \sin^2 \theta_i}} \right|^2. \end{aligned} \quad (2.20)$$

A further common assumption in computer graphics is that light is unpolarized, i.e. equally polarized in parallel and perpendicular directions. Hence, the

reflectance is calculated by averaging the parallel and perpendicular reflectance:

$$R = 0.5(R_{\parallel} + R_{\perp}). \quad (2.21)$$

2.1.4 Reflection functions

The material reflection behaviour is crucial to predict the light propagation within a scene. A complex function of 16 variables is required to account for all physical phenomena. This function is called general reflection function (GRF). The $GRF(\lambda_i, x_i, y_i, z_i, t_i, \theta_i, \phi_i, \lambda_o, x_o, y_o, z_o, t_o, \theta_o, \phi_o, \theta_t, \phi_t)$ describes the illumination at the location x_i, y_i, z_i by the wavelength λ_i originating from the spherical incident angles $\omega_i(\theta_i, \phi_i)$ in time t_i and the observation of the wavelength λ_o under the spherical reflection angles $\omega_o(\theta_o, \phi_o)$ at the surface location x_o, y_o, z_o in time t_o . The corresponding spherical transmission angles are defined by $\omega_t(\theta_t, \phi_t)$. It is currently infeasible to measure or build a mathematical model of such a complex function. Hence simplifying assumptions are inevitable to use reflection functions in practical applications.

Taxonomy

Referring to the taxonomy in Figure 2.1, all mentioned reflection models, excluding dynamic textures, are derived from the bidirectional surface scattering reflectance distribution function (BSSRDF).

The BSSRDF was proposed by Nicodemus et al. [14] in 1977. It is a nine-dimensional model describing the scattering behaviour of materials with subsurface scattering characteristics. This model takes into account not only the spherical incident (θ_i, ϕ_i) and reflection angles (θ_o, ϕ_o) , but also the illumination x_i, y_i and reflection surface location x_r, y_r . The BSSRDF is an accurate but still complex scattering model making measurements very challenging. Until now, no satisfactory BSSRDF data has been acquired [1]. Therefore, further simplified models are used in computer graphics, which can be divided into textured and homogeneous materials. Textured materials have surface variations on the macroscopic scale that can be seen with the naked eye. Spatially varying models are necessary to properly describe textured materials' reflection since the reflection may change from point to point. Homogeneous materials have surface variations on the microscopic scale, which is why they are perceived as homogeneous. These materials need only a point-wise description, i.e. the same function is evaluated at different surface points.

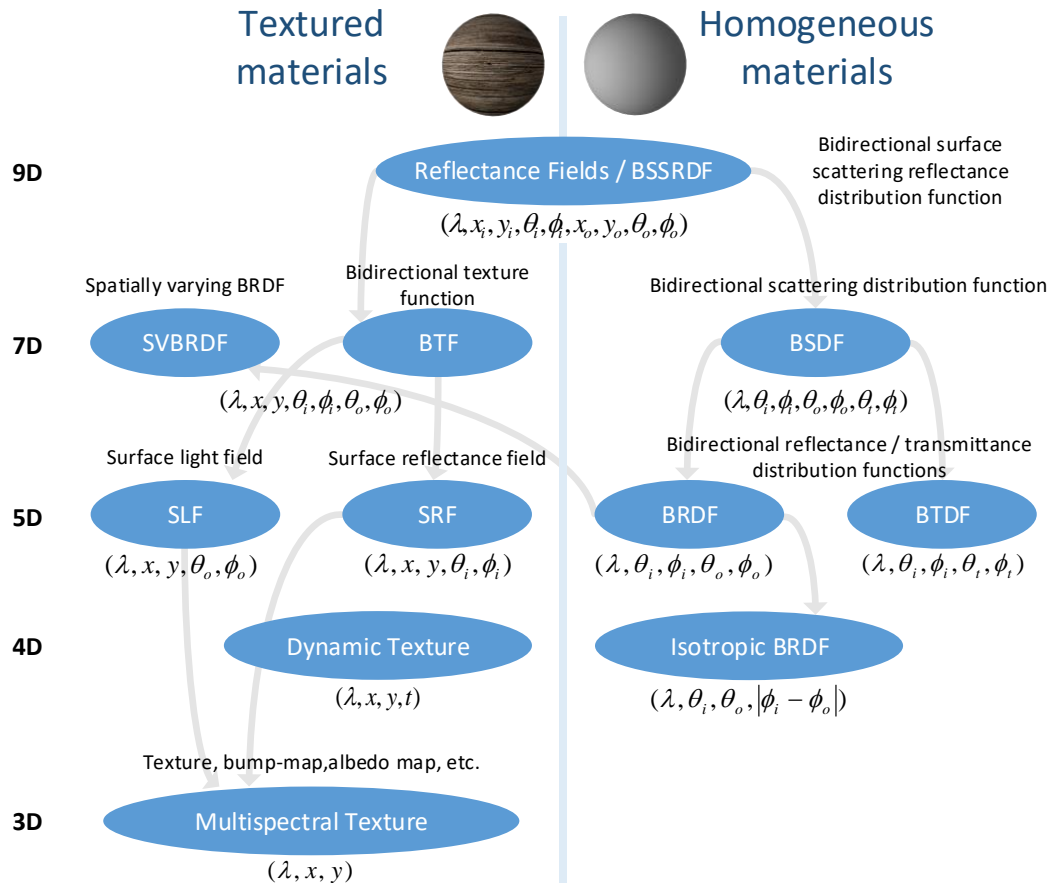


Figure 2.1: A taxonomy of selected reflectance models (According to [1])

BRDF

The bidirectional scattering distribution function (BSDF) approximates the light-interaction with homogeneous materials (see Figure 2.1). The BSDF comprises the reflection of light on the surface and the transmission of light through the material. The two scattering effects are considered individually in the bidirectional reflectance distribution function (BRDF) and the bidirectional transmittance distribution function (BTDF). However, the BRDF gets more attention in computer graphics, probably because most real-world materials are opaque.

The BRDF describes how an incident light beam from a particular direction is scattered at a surface within the upper hemisphere. It is a five-dimensional function that returns the amount of light of a certain wavelength coming from and going in a certain direction. Figure 2.2 shows a schematic diagram of a surface area dA illuminated from the direction $\omega_i(\theta_i, \phi_i)$ by a radiation cone within the solid angle $d\Omega_i$ and its reflection in the direction $\omega_o(\theta_o, \phi_o)$ within the solid angle $d\Omega_o$. The BRDF is defined as the ratio of the reflected radiance to the incident




Figure 2.2: The BRDF is defined as the reflected spectral radiance dL_o within the cone $d\Omega_o$ towards ω_o divided by the incident spectral irradiance E_i within the cone $d\Omega_i$ coming from ω_i . (According to [2])

irradiance, as shown in the following equation:

$$f_r(\lambda, \omega_i, \omega_o) = \frac{dL_o(\lambda, \omega_o)}{dE_i(\lambda, \omega_i)}, \quad (2.22)$$

where $dL_o(\lambda, \omega_o)$ is the reflected radiance in the direction ω_o caused by the irradiance $dE_i(\lambda, \omega_i)$ from the direction ω_i .

This thesis primarily deals with the so-called in-plane BRDF. The in-plane BRDF describes a three-dimensional subspace of the five-dimensional BRDF space, where only incident and reflection directions located in the same plane are evaluated. Hence, the spherical incident and reflection directions reduce to the polar angles θ_i and θ_o ranging from -90° to 90° , where 0° is normal to the surface.

BRDF representations

There are different ways to describe the BRDF of a material. In Figure 2.3 the different BRDF representations are classified in a continuum diagram. The left end represents BRDFs that are physically correct but not intuitive and hard to manipulate. On the contrary, the right end represents BRDFs that are not physically correct but intuitive and easy to manipulate. From left to right, the BRDFs are grouped into measured, data-driven, physically based, and empirical BRDF models. The measured BRDF consists of several optical measurements, the data-driven models include a numerical process, and the empirical and physically based models are analytic formulations.

This thesis deals with measured BRDFs and physically based BRDF models. Measured BRDFs are essential for validating and developing new BRDF models; physically-based BRDFs, on the other hand, provide a good trade-off between physically correct and intuitive. In the following, both BRDF representations are explained in more detail.




Figure 2.3: Classification of BRDF representations in a continuum diagram.

BRDF measurement: The BRDF measurement is complex since it is a five-dimensional function. It requires a device consisting of a light source illuminating the material sample, a detector measuring the reflected light, and a sample holder. Usually, two different setups are used: an image-based or a gonioreflectometer-based setup.

The image-based setup performs measurements by taking an image of a curved or spherical sample. Due to the sample geometry, one image covers many combinations of incident and reflection directions. Consequently, the image-based setup drastically shortens acquisition time but only provides RGB values on the flip side. The approach was first introduced by Lu et al. [15] but limited to in-plane BRDFs. Marschner et al. [16] improved the idea by measuring the isotropic BRDF. In 2003, Matusik et al. [17] used a similar approach to acquire dense BRDFs of more than 130 materials (MERL database). The MERL database has been widely used to develop and validate reflection models. However, the limitation on RGB values and measurement errors caused by lens aberrations [18] make it difficult to develop more sophisticated BRDF models, including complex physical phenomena such as diffraction.

The gonioreflectometer approach, as described in [2, 19–24], has more mechanical degrees of freedom compared to the image-based approach. Usually, one of the three components, detector, light source, or sample, is fixed, and the other components are flexible to measure the BRDF in-plane or in the hemisphere. Instead of cameras, spectrometers are often used, which sample the reflection spectrum densely. Along with high angular accuracy, these are the main advantages of the gonioreflectometer-based setup. However, sampling densely the hemisphere is time-consuming because each combination of incident and reflection direction must be addressed individually. Smart sampling strategies considerably speed up

the BRDF acquisition. Recently, Dupuy et al. [24] presented a dynamic sampling strategy, where the reflection behaviour is first estimated by back-scattering measurements and then sampled corresponding to the estimated information density. This sampling strategy heavily reduces the measurement time but is still significantly slower than the image-based setup.

Physically based BRDF models: As aforementioned, measuring the BRDF is time and cost-intensive and special equipment is required. Furthermore, measured BRDFs don't support importance sampling, which complicates their application in rendering. Therefore, analytic BRDFs, mainly physically based BRDFs, are widely used in computer graphics.

As follows from its name, physically based BRDFs are based on physical parameters and model the material reflectance in a physically plausible manner. Usually, they obey positivity, Helmholtz reciprocity, and energy conservation [25]:

1. Positivity: $f_r(\lambda, \omega_i, \omega_o) \geq 0$,
2. Helmholtz reciprocity: $f_r(\lambda, \omega_i, \omega_o) = f_r(\lambda, \omega_o, \omega_i)$,
3. Energy conservation: $\int_{\Omega} f_r(\lambda, \omega_i, \omega_o) \cos\theta_i d\omega_i \leq 1$.

Positivity ensures that the material cannot absorb more energy than it receives, Helmholtz reciprocity means that the incident and reflected direction can be reverted without affecting the BRDF, and energy conservation guarantees that the material does not emit energy.

Physically based BRDF models can be divided into models based on geometric- and wave-optics. As explained in Section 2.1, geometric optics considers the energy propagation along light rays. The underlying assumptions heavily simplify the computations, consequently allowing for the usage in real-time applications. Due to their simplicity, geometric optics models are much more popular than wave-optics models. However, to simulate physical phenomena such as diffraction or interference, wave-optics have to be taken into account.

a) Geometric optics models: Most physically based BRDF models use the microfacet theory to describe the light reflection on rough surfaces. The microfacet theory approximates the surface's geometry by a collection of tiny reflectors (microfacets) that are randomly distributed and usually assumed as mirrors. It was first proposed by Bouguer [26] and further developed by Beckmann and Spizzichino [27], as well as by Torrance and Sparrow [28]. In 1982, Cook and Tor-

rance [29] introduced a BRDF model based on microfacet theory to the computer graphics community, which is still one of the most popular BRDF models.

The Cook-Torrance model describes the reflection of light from a rough material as a combination of a diffuse and a specular component. The diffuse component represents both light that is refracted and randomly re-emitted and light that is reflected multiple times on the surfaces. This component is assumed as Lambertian; hence, it is independent of the incident direction. The specular component represents the first reflection on the surface and is approximated by a microfacet model.

The Cook-Torrance model is defined as:

$$f_r(\lambda, \omega_o, \omega_i) = \frac{k_{d,\lambda}}{\pi} + k_{s,\lambda} \frac{D(\omega_h)F(\omega_i, \omega_h)G(\omega_o, \omega_i, \omega_h)}{4\cos(\theta_o)\cos(\theta_i)}. \quad (2.23)$$

The diffuse component is described by a wavelength-dependent diffuse albedo $k_{d,\lambda}$ normalized by Pi. The specular component consists of a wavelength dependent specular albedo $k_{s,\lambda}$, the normal distribution function $D(\omega_h)$, the Fresnel function $F(\omega_i, \omega_h)$ and the shadowing-masking function $G(\omega_o, \omega_i, \omega_h)$, as well as a normalization factor.

The normal distribution function returns the differential area of microfacets, whose normals are aligned with the half-direction ω_h . Only these microfacets contribute to the scattering due to the mirror assumption. The shadowing-masking function returns the fraction of microfacets that are not masked or shadowed. And the Fresnel function defines the reflectance of each microfacet, which is precisely described by the Fresnel equations (Eq. 2.19). We refer the reader to the paper of Cook and Torrance [29] for the particular implementation of the model functions.

In the last decades, several implementations of the normal distribution function, shadowing-masking function, and Fresnel function were introduced [18, 29–37]. In this thesis, we use the GGX distribution function [34] with the associated Smith shadowing-masking function [38]. Smith approximates the G function as the separable product of two monodirectional masking functions G_1 :

$$G(\omega_o, \omega_i, \omega_h) \approx G_1(\omega_o, \omega_h)G_1(\omega_i, \omega_h). \quad (2.24)$$

The GGX distribution function and Smith's masking function are defined as:

$$D_{GGX}(\omega_h) = \frac{\alpha_g^2}{\pi \cos^4 \theta_h (\alpha_g^2 + \tan^2 \theta_h)^2}, \quad (2.25)$$

$$G_1(\omega_o, \omega_h) = \frac{2}{1 + \sqrt{1 + \alpha_g^2 \tan^2 \theta_o}}, \quad (2.26)$$

where α_g is the surface roughness.

We use different Fresnel implementations for dielectrics and conductors. The exact Fresnel equations (Eq. 2.19) with the refractive index of aluminium are used for the conductive aluminium samples. For the dielectric ColorChecker patches, the Fresnel approximation by Cook and Torrance is used to simplify the fitting against measured in-plane BRDFs. Under the assumption of unpolarized light, Cook and Torrance approximate the Fresnel equation as follows:

$$F_{Cook-Torrance}(\omega_i, \omega_h) = \frac{1}{2} \frac{(g-c)^2}{(g+c)^2} \left\{ 1 + \frac{[c(g+1)-1]^2}{[c(g-c)+1]^2} \right\} \quad (2.27)$$

with $g = \sqrt{\eta - 1 + c^2}$ and $c = \cos(\beta)$, where β is the angle between ω_h and ω_i .

b) Wave-optics models: Two historical approaches analytically describe light-matter interaction: the Rayleigh–Rice (RR) vector perturbation theory [39, 40] and the Beckmann–Kirchhoff (BK) scattering theory [41]. Both are complementary in their range of validity. On the one hand, the Rayleigh-Rice theory is valid for smooth surfaces for arbitrary incident and reflection angles. On the other hand, the Beckmann-Kirchhoff theory is valid for rougher surfaces but only for small incident and reflection angles. The Generalized Harvey Shack (GHS) theory [42] combines the advantages of the previous approaches without their limitations.

A few BRDF models were developed based on these theories. He et al. [43] extended the Cook-Torrance model using the BK scattering theory. Their model matches well the measured data of a few materials. Nevertheless, Ngan et al. [44] showed that when fitting against the MERL database, He’s model leads to similar results as the Cook-Torrance model.

Löw et al. [36] introduced a BRDF model based on the RR theory. Even though it is limited to smooth surfaces, it represents such cases from the MERL database very well. A different approach is presented by Dong et al. [45]. For the first time, the anisotropic appearance of metal samples was determined only by their microgeometries. They demonstrate that both geometric-based and Kirchoff-based approaches can model the appearance properly. However, a systematical validation of the proposed model with measured data is missing.

Holzschuch and Pacanowski [46] derive the surface reflectance by a superposition of the reflectance on the macro- and microgeometry. The first is characterized by the standard Cook Torrance model and the second by the modified

Harvey Shack theory. The model can reproduce some wavelength-dependent phenomena observed in the MERL database but still carries the limitations of the modified Harvey Shack approach.

Yan et al. [3] derived a generalized BRDF representation where the three scalar diffraction theories, Harvey-Shack, Generalized Harvey Shack, and Kirchhoff, can be alternatively used. Their model allows the simulation of full diffraction effects of arbitrary micron-scale heightfield geometries.

2.2 Gonioreflectometer

Ideally, a gonioreflectometer captures the reflection behaviour densely in spatial, spectral, and angle domains. In practice, each device addresses only one or two of these domains. Our custom-built gonioreflectometer focus on the spectral domain, thus capturing highly resolved reflectance spectra at a single spot. Due to the spot measurement, our gonioreflectometer is restricted to isotropic, homogeneous, and planar samples.

2.2.1 Setup

Our gonioreflectometer, as illustrated in Figure 2.4a, consists of three components: a collimating light system, a focusing detector combined with a high-precision spectrometer, and a two-axis gonio-stage. In the following, the individual components are explained in more detail and two setup configurations are presented.


Light system

The light system consists of the Schott KL2500LCD¹ light source connected with a projector optic. The light source uses a halogen radiator with a continuous spectrum in the visible wavelength range, which must be warmed up for at least 30 minutes before measurements to avoid intensities fluctuations. The projector is mounted on the optical table and provides a collimated and uniform illumination spot in the centre of the sample.


Detector

The detector consists of a detector head and a spectrometer. The detector head is a lens tube including an achromatic lens, which focuses the reflected light from the sample onto a fibre cable connected with the high-sensitivity spectrometer


¹http://www.schott.com/lightingimaging/english/microscopy/products/kl/2500lcd.html?highlighted_text=kl+2500+lcd



(a) Gonioreflectometer setup.



(b) Light system.



(c) Detector.

Figure 2.4: Our custom-built gonioreflectometer (a) consists of a light system (b) producing a uniform and collimated illumination, a detector (c) combining a focusing detector head with a high-precision spectrometer, and a two-axis goniostage rotating the sample and detector head.

QE Pro by Ocean Optics². The spectrometer provides a high dynamic range of $\sim 85000:1$ and measures in a wavelength range of 250 to 1050 nm in 0.8 nm intervals. The sensor temperature is controlled by a thermoelectric cooling (TEC) device, considerably reducing the thermal noise. Hence the spectrometer has an excellent signal-to-noise ratio (SNR) of 1000:1. The detector's measurement spot has a diameter of around 2.5 mm when measuring normal to the sample. At measurement angles deviating from the normal direction, the measurement spot is horizontally enlarged by one over the cosine of the measurement angle. The maximum measurement angle is 85° ; consequently, the maximum measurement spot has a horizontal diameter of approximately 30 mm. Since, in our setup, the detector defines the area of the BRDF quantity, the maximum measurement spot must be within the illumination spot and sample size.

Gonio-stage

The last component of the gonioreflectometer is the gonio-stage. In our setup, the gonio-stage consists of two rotary measuring stages by OWIS³. The two stages enable the independent rotation of the sample holder and detector head around axis 1 and 2, as shown in Figure 2.4a, with repeatability of $< 0.01^\circ$.

Configurations

In this thesis, we use two gonioreflectometer configurations: the first for diffuse dielectrics and the second for conductors with a strong specular reflection that is not well captured by the first configuration. The first configuration is shown in Figure 2.4 (a). As illustrated, the sample holder is tilted by 4° and the detector head by 8° . In this way, the detector does not shadow the illumination spot. Consequently, back-scattering can be acquired more precisely. In the second configuration, the detector and light source are always on a plane perpendicular to the sample. Hence, the illumination and measurement directions are in-plane. Although only the second configuration measures the in-plane BRDFs properly, we will refer in both cases to in-plane BRDF.

2.2.2 Calibration

The spectrometer of our gonioreflectometer natively counts electrons that are emitted when light hits the photosensitive sensor. We use the relative calibration method to calculate the BRDF from the electron counts. In the following, we

²<http://oceanoptics.com/product/qe-pro/>

³http://www.owis.eu/fileadmin/user_upload/owis.eu/products/pdf/pi_dmt_100.pdf

briefly review the relative calibration method as described by Sing Foo in his master thesis [2].

When measuring the BRDF f_r (Eq. 2.22) with a gonioreflectometer, it is evaluated at discrete values and the solid angles of the detector and light source are finite. To symbolize that the differential quantities are, for now, discrete, we refer to them by the Δ symbol.

The source irradiance ΔE_i received by the illuminated sample can be defined as:

$$\Delta E_i(\omega_i) = \Delta E_0 \cos(\theta_i), \quad (2.28)$$

where ΔE_0 is the source irradiance in the normal direction. The reflected radiance $\Delta L_o(\omega_i, \omega_o)$ is proportional to the measured signal $V(\omega_i, \omega_o)$ and thus, equal to the measured signal divided by a calibration factor C_λ .

$$\Delta L_o(\omega_i, \omega_o) \propto V(\omega_i, \omega_o) \quad (2.29)$$

$$\Delta L_o(\omega_i, \omega_o) = V(\omega_i, \omega_o) C_\lambda^{-1} \quad (2.30)$$

In terms of Equation 2.28 and Equation 2.30, f_r can be written as:

$$\begin{aligned} f_r(\omega_i, \omega_o) &= \frac{\Delta L_o(\omega_i, \omega_o)}{\Delta E_i(\omega_i)}, \\ &= \frac{V(\omega_i, \omega_o)}{C_\lambda \Delta E_0 \cos(\theta_i)}, \end{aligned} \quad (2.31)$$

where for simplicity, the wavelength dependence of the f_r is left out. To indicate that the calibration factor C_λ is computed for all wavelengths, it is subscripted by λ . The unknown term $C_\lambda \Delta E_0$ is independent of the incident and reflection direction. It can be determined by measuring the reflection of the diffuse standard ($V_{standard}$) and using the associated BRDF value ($f_{standard}$) in the following way:

$$C_\lambda \Delta E_0 = \frac{V_{standard}(\omega_i, \omega_o)}{f_{standard}(\omega_i, \omega_o) \cos(\theta_i)}. \quad (2.32)$$

In this thesis, we use the reflectance standard Spectralon by Labsphere⁴ together with the BRDF values provided by Durell et al. [47]. Durell et al. conducted a round-robin test with four laboratories, each measuring the BRDF of two Spectralon samples. The averaged BRDF of two laboratories were denoted as reference due to their great accordance.

Although the calibration factor C_λ is wavelength dependent, we determine it only for the wavelength 533 nm. This is a reasonable assumption due to the small

⁴<https://www.labsphere.com/product/spectralon-diffuse-reflectance-standards/>

reflectance variance of Spectralon in the visible spectrum between 360 - 740 nm, which is $<0.5\%$ according to the technical guide of Spectralon. The calibration factor is calculated for several measurement angles and then averaged to minimize the influence of measurement errors. More precisely, C_λ is calculated for the incident angle $\theta_i = 10^\circ$ and the reflection angles $\theta_o = -80^\circ, -70^\circ, \dots, 70^\circ, 80^\circ$.

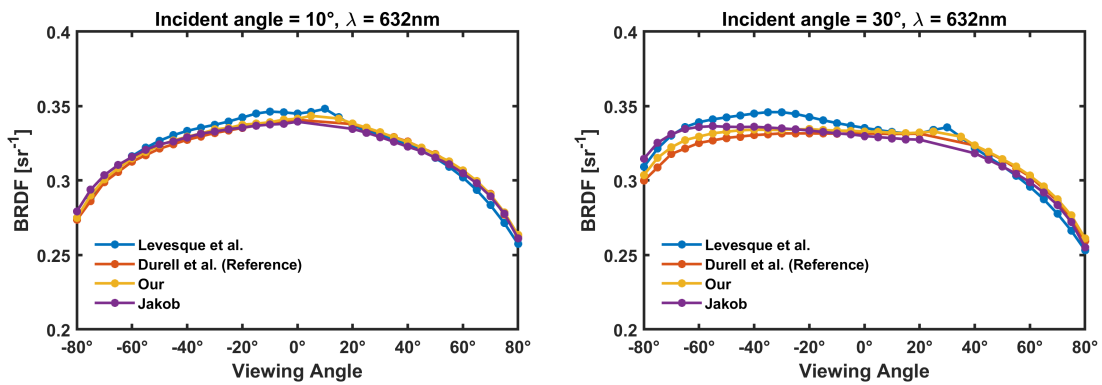
2.2.3 Cross-validation

As stated in the previous section, a gonireflectometer is a complex measurement instrument consisting of several components. Each component is a potential source of error; thus, BRDF measurements are error-prone. To evaluate the fidelity of our gonireflectometer, we conducted cross-validation by comparing our measured Spectralon BRDF with the reference data by Durell et al. [47], as well as with data provided by Levesque et al. [48, 49] and Wenzel Jakob [50]. The results of the cross-validation were presented on the Workshop on Material Appearance Modeling [51].

The cross-validation is divided into two parts. First, we compare the measured BRDFs at small incident angles and a single wavelength to provide information about the fidelity of Levesque's, Jakob's and our gonireflectometer. Second, we compare the BRDF spectra at grazing incident angles to investigate an unusual wavelength shift we observed at Spectralon under grazing angles. This effect was first documented by Levesque et al. [48, 49], but they did not provide a reasonable explanation. Thus, measurement errors can not be excluded as a source.

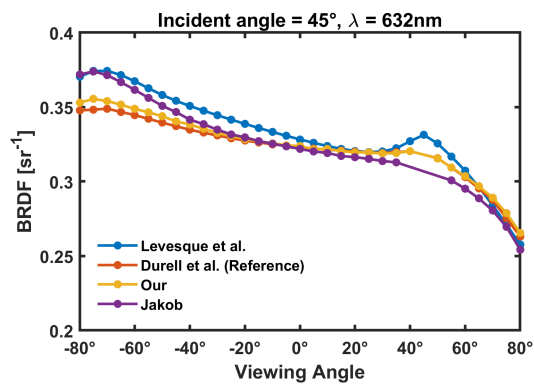
In Figure 2.5, Levesque's, Jakob's, and our BRDF values at 632 nm are compared to the reference data by Durell et al.. The plots (a), (b), and (c) illustrate the measured data at the incident angles of 10° , 30° , and 45° , respectively. The figure shows that our and the reference BRDF match very well. The small average and maximum error of 0.7% and 2.5% confirm the high fidelity of our gonireflectometer. Levesque's and Jakob's measurements match the reference data for small incidence angles, while with increasing incident angles, deviations in the forward scattering can be observed. These deviations probably originate from different Spectralon samples used by each laboratory. We suspect Levesque et al. used a slightly smoother Spectralon sample due to the higher specular reflection.

The reference data by Durell et al. are only available for small incident angles and only for two wavelengths. Therefore, we exclude them from the spectral comparison in Figure 2.6. The figure depicts Levesque's, Jakob's, and our BRDF spectra at the incident angle of 80° and the reflection angles from -50° to -80° .



(a) BRDF at: incident angle 10°.

(b) incident angle 30°.



(c) incident angle 45°.

Figure 2.5: Spectralon BRDF at 632 nm and the incident angles 10°, 30° and 45° acquired by Levesque et al., Durell et al., Jakob and us.

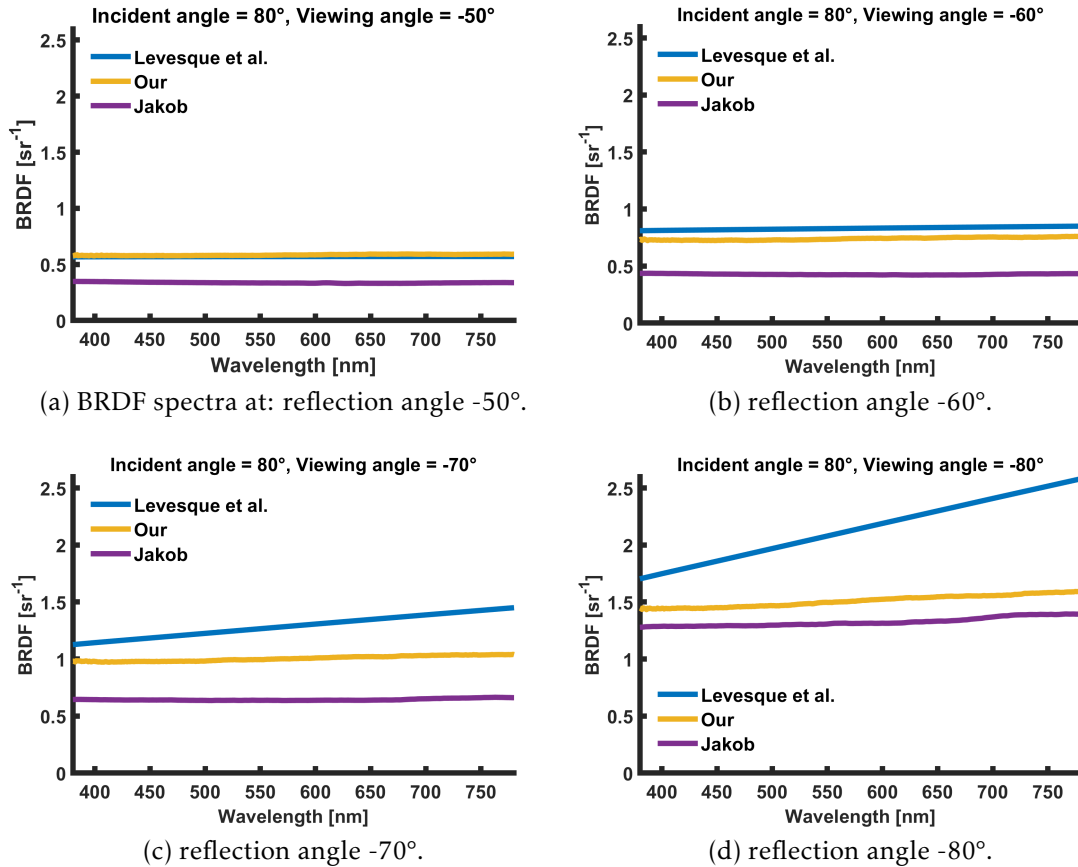


Figure 2.6: Spectral BRDF of Spectralon at an incident angle of 80° and the reflection angles -50° , -60° , -70° and -80° acquired by Levesque et al., us and Jakob.

Levesque’s and our spectra match well for small reflection angles but deviate for large ones. The slope of Levesque’s spectra, i.e. the wavelength shift, is much stronger. These deviations are consistent with the assumption that Levesque et al. probably used a slightly smoother Spectralon sample. Jakob’s spectra are consistently lower than ours while having a similar slope. We don’t have a meaningful explanation for Jakob’s deviations. Perhaps, they are caused by the different measurement setups.

Despite the deviations, all three sources show a tilting of the reflectance spectra (wavelength shift), especially at large reflection angles. Each of the three gonioreflectometers has an individual setup with differing components, but all use a detector with a focusing achromatic lens. Although we can rule out the measurement errors as the source of the wavelength shift, we have not yet proven that the optical aberration is not the actual source.

2.2.4 Detector validation

To exclude the achromatic lens as the source of error, we conducted another Spectralon measurement with a modified gonireflectometer setup. In the modified setup, the achromatic lens is replaced by an integrating sphere and the illumination is converged instead of collimated. In Figure 2.7 (a) and (c) both gonireflectometer setups are illustrated. Note that our gonireflectometer is not designed for the modified setup, which is why the modified setup has a considerable lower angular resolution than the original setup. Furthermore, for the modified setup, a bigger Spectralon sample is necessary. Hence, we used a larger white reflectance standard with a slightly smoother surface than the original Spectralon sample.

Figure 2.7 depicts the BRDF spectra of the diffuse standard for the incident angle of -80° and the reflection angles from 0° to 80° acquired with the original (b) and modified setup (d). There are clear differences between the two setups. The modified setup has much higher BRDF spectra than the original setup, especially at small reflection angles. These deviations can be traced to the different angular resolutions of the setups. Nevertheless, the shape of the spectra remains similar, and both setups have the same wavelength shift. This observation confirms that the wavelength shift is a property of the material reflectance function and discards the lens aberration hypothesis.

2.3 Other measurement devices


In addition to the gonireflectometer, we used other measuring instruments in this work, which are briefly presented below.

Konica Minolta - CS-1000A

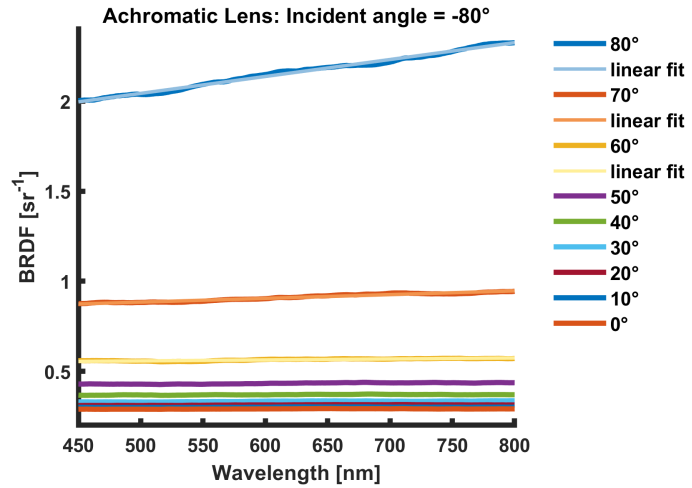
The CS-1000A by Konica Minolta is a portable high-precision spectroradiometer. It accurately measures the spectral radiance in a wavelength range from 380 to 780 nm in 0.9 nm intervals within an acceptance angle of 1 degree. It provides a radiometric accuracy of 2 % in a luminance range of 1 to 8000 cd/m^2 , and a spectral accuracy of 0.3 nm. The CS-1000A is calibrated on the standard illuminant A and we used it with the 50 nm standard lens.

X-Rite - EyeOne Pro 2


The EyeOne Pro 2 is a handy spectroradiometer by X-Rite, usually used to calibrate displays, printers, or projectors. The device supports a contact and distance mode. The former enables the measurement of the reflected or emitted light. In



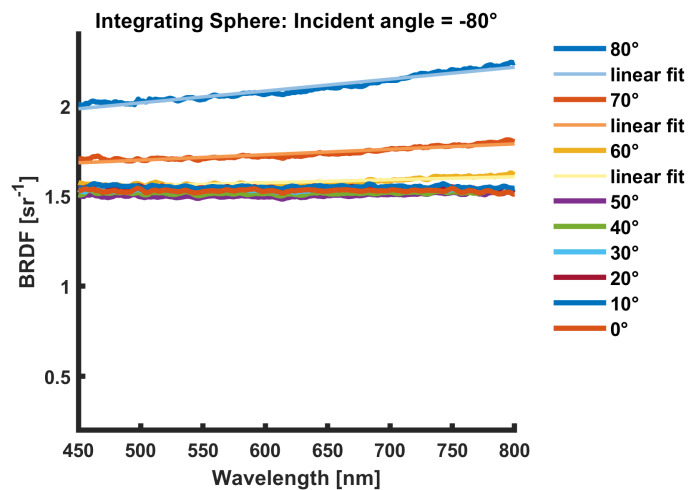
(a) First: Detector with achromatic lens.



(b) BRDF spectra acquired with first detector.



(c) Second: Detector with integrating sphere.



(d) BRDF spectra acquired with second detector.

Figure 2.7: In the first detector setup (a), the sample is illuminated by collimated light. The scattered light is captured and coupled into a fibre cable by an achromatic lens. The fibre cable is connected to the spectrometer QE Pro. In the second detector setup (c), the achromatic lens is replaced by an integrating sphere from Thorlabs and the sample is illuminated by converged light. In both detector setups, the BRDF spectra include the linear wavelength shift at grazing angles (b,d).

both cases, the ambient light is excluded from the measurement. The spectroradiometer measures the spectral radiance in a wavelength range of 380 to 730 nm in 10 nm intervals. The inter-instrument-agreement is in average $0.4\Delta E_{94}^*$ and maximal $1.0\Delta E_{94}^*$. The manual doesn't further specify the measurement error.

Macro photos

We built a simple goniometer that captures macro photos under a fixed camera position and varying illumination angles to obtain spatial information on the light interaction with rough surfaces. The goniometer consists of a collimated high-power LED, which can be moved on a semi-circle around the sample, and a Nikon D810 with a micro Nikkor 105 mm lens mounted on a tripod at 45° in regards to the sample surface. The camera is calibrated by performing a manual white balance with a Spectralon sample illuminated with the LED light source.

Confocal microscope

The NanoFocus μ surf Generation C is a white light confocal microscope. It consists of a compact confocal measuring head, which is arranged on a stable work table so that it can be moved along the z-axis by a motor. The sample is placed on a computer-controlled measuring table. When performing a non-contact measurement of the sample topography, the sample is first positioned on the measuring table and then sampled on the z-axis by gradually moving the measuring head. In combination with the 320S lens, the NanoFocus μ surf has a lateral and vertical resolution of 625 nm and 20-40 nm, respectively. A single shot captures an area of $320 \times 320 \mu\text{m}$ depending on the vertical resolution. A stitching mode enables the measurement of more extensive areas up to 7.3 mm^2 .

Virtual gonioreflectometer

The virtual gonioreflectometer was implemented by Yang Chen, who is a PhD student at TU Delft under the supervision of Ricardo Marroquim. The tool decomposes the surface scattering on rough surfaces into individual reflection paths, which helps us to better understand observed scattering phenomena on rough surfaces. For the decomposition, the measured microgeometry and refraction index are required. The simulation is based on geometric optics and assumes each facet of the triangulated measured surface as an ideal mirror; thus, the Fresnel equations describe the facet's reflectance. In the light simulation, the surface is virtually radiated with a statistically critical number of rays traced on the surface. This gathers information about the energy coming from rays that reach the sensor after a specific number of reflections.

3. Validation of light simulation

Physically based rendering has become a popular technique for offline and real-time rendering and has been adopted in many industrial applications. It is based on physical principles to simulate light propagation and -scattering (light simulation), and its primary goal is usually realism [52, p.1]. The tremendous advancements of physically based rendering have led to photorealistic images but often fail to be predictive [53, p.11]. Deviations of rendered images from reality are not critical in the entertainment segment but considerably limit the benefits of VP, as stated in Chapter 1.

In physics, the light simulation is described in different levels of accuracy, forming a hierarchy of increasing approximations. Quantum electrodynamics (QED) is the most comprehensive theory describing all phenomena involving electrically charged particles interacting through the exchange of photons. A simplified approach is wave-optics describing light scattering by means of electromagnetic fields and Maxwell's equations, thus describing wave phenomena such as diffraction and interference. The bottom of the hierarchy is geometric optics that do not consider the wave nature of light.[54]

Computer graphics usually further simplify the geometric optics approach by neglecting polarisation, assuming the refractive index as piecewise constant and approximating the electromagnetic spectrum by RGB-values. At the same time, it is assumed that the physical deviations are not perceptible due to the limitations of the human eyes. Only a few works verify the made assumptions experimentally, which makes it difficult to assess how well current render algorithms predict reality.

A common way to determine the influence of these approximations on the appearance of real materials in a natural lighting environment is the comparison of rendered images with reference data captured in a real-world scene. This comparison method is called the experimental verification method and requires the physical description of the scene, as well as accurate reference data. In the following, we summarise both and refer to them as ground truth data (GTD). An important observation is that the fidelity of the verification method strongly depends on the accuracy of the GTD; hence its validation is necessary.

In this chapter, we introduce a new reference scene (Normbox) similar to the Cornell box, which we used to acquire an accurate and spectrally resolved set of GTD. It consists of the description of the Normbox, including isotropic BRDFs of 24 colour patches and the reference measurements of all patches under 13 different angles inside the Normbox. Our reference data covers rough materials with varying reflectance spectra and illumination, from direct to indirect light dominated.

We register all measurement errors during the acquisition but estimating their influence on the reference targets within the box is challenging. Therefore, we additionally implement the experimental verification method, using the reference renderer Mitsuba [55] to predict the spectral radiance of the target patches. We determine the physical and perceived differences between the reference and predicted data to assess the fidelity of the GTD. Hereby, is the high resolved spectral radiance as reference data a key advantage compared to previous approaches, which only use colourimetric quantities and single light measurements.

3.1 Related work

Greenberg et al. [56] proposed a framework for verifying photorealistic renderers. This framework is divided into three steps: light reflection model, global light transport simulation, and image display. According to the authors, each step must be verified individually by comparing it with the real world. The first two steps may be verified by comparing physical parameters, while the third step requires a perceptual comparison.

This chapter deals with the first two steps. We use the experimental verification method to investigate the influence of the approximated light reflection and global light transport on the appearance of real-world materials. In the following, we present related work on the experimental verification method.

In 1984, Goral et al. [57] introduced the experimental verification method to the computer graphics community. They presented the well-known Cornell Box, a cube of fiberboard panels painted with diffuse latex paints. Three sides are painted white, while the left and right sides are painted red and blue to investigate the effect of colour bleeding. The remaining side of the cube is open to take pictures and illuminate the scene with diffuse light. In their work, they did not perform an accurate comparison to quantify the physical or perceptual differences between the simulated and real Cornell Box. Their approach only allows them to state that the colour-bleeding effect is visible in both cases.

Meyer et al. [58] followed up on the Cornell Box approach. They slightly changed the Cornell Box (e.g. they mounted a diffuse light source in the ceil-

ing of the cube and placed two small boxes within the cube) and extended the approach with a radiometric comparison. They took irradiance measurements at 25 locations within the reference scene and compared them with rendered images using a radiosity method. The resulting differences are below 4%. A drawback of their approach is the irradiance measurements' limited spatial and spectral resolution. It is difficult to accurately identify the measured points in the rendered image, and due to the missing spectral resolution, it is impossible to evaluate the colour differences.

Greenberg et al. [56] presented a framework for realistic image synthesis. The framework extends the work of Meyer et al. by increasing the spatial and spectral resolution of the reference data. They acquired spectral images using a camera with a 1280 x 1024 CCD array and eight bandpass filters in front of the camera. However, they did not provide results of the described framework. Pattanaik et al. [59] implemented the presented framework but only provided a single reference and predicted image, as well as the resulting difference image, without communicating more detailed results. However, the given reference and predicted image already show clearly visible differences.

Accurate verification of the global light transport simulation was conducted by Schregle and Wienold [60]. They broke the verification process into simple component studies that can be checked individually and combined to build more complex compound case studies. A simple setup consisting of a box and an area light source was used as a reference scene. Diffuse patches or a brushed aluminium plate were integrated depending on the study case. Flexible sensors in the ceiling measured the direct illuminance at different positions, representing the reference data. The authors compared the reference data to a forward and backward raytracing solution, namely Photon Map [61] and Radiance [62, 63]. The difference between the reference and predicted illuminance is in all conducted case studies on average below 3%.

A more recent work by Bärz et al. [64] followed a similar approach as Meyer et al. [58]. Instead of a light source on the cube's ceiling, they placed a monitor on the left side. On the opposite side, they mounted a Macbeth ColorChecker and measured all patches of the ColorChecker with a colourimeter as reference data. They conducted two perceptual comparisons; the comparison of the measured reference data with, first, the predicted data itself and second, with the predicted data displayed on a monitor, identical to the used light source. Both comparisons achieve colour differences below the JND for most patches. However, there are also several patches with extraordinarily high colour differences.

The works described so far use a reference scene similar to the Cornell Box, which provides a simple and constant environment. Nonetheless, natural scenes

consist of more complex geometries, materials, and light conditions, which cannot be validated with these simple setups. Some works proposed more realistic reference scenes like a car, a dashboard, an office, or an atrium [65–69]. These works compare the rendered images with photos or several luminance measurements to determine their fidelity. The major problem of these approaches is that it is impossible to accurately determine the GTD of such complex scenes. In particular, scenes, where sunlight is present, are useless for an accurate comparison since the illumination conditions are constantly changing.

3.2 Validation framework

Similar to the related works, our validation framework is based on the experimental verification method. As illustrated in Figure 3.1, the framework is divided into two paths: the comparison path (yellow) and the error path (blue).

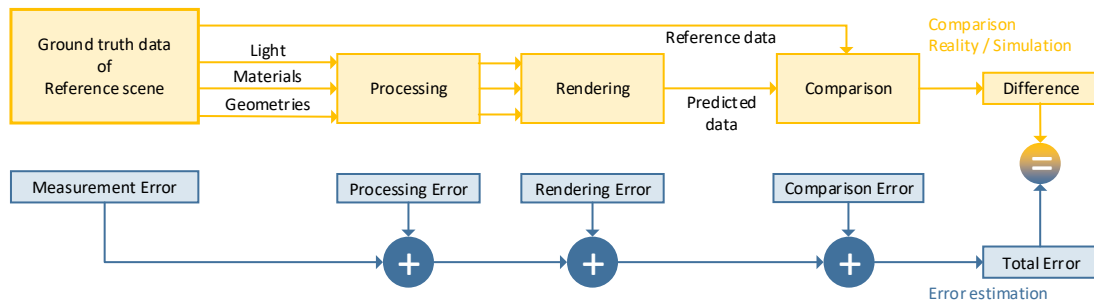


Figure 3.1: Overview of our validation framework.

The comparison path comprises the experimental verification method divided into four steps. In its first step, the GTD of the reference scene is acquired, consisting of the light emission- and material reflection distributions, scene geometries, and reference data. In the second step, this data is processed for the rendering phase, in which the reference scene is simulated (predicted data). Finally, in the comparison step, the predicted and reference data are compared to determine the deviations from reality. The resulting differences are composed of various errors considered in the error path.

The error path describes all errors introduced at every step of the comparison path. Each error contributes to the total error, which corresponds to the determined difference in the comparison path. In practice, estimating the total error is challenging or even impossible because many errors contribute in a complex way. Nevertheless, the individual errors help us to indicate the primary error sources and give us information about the GTD fidelity.

3.3 Implementation

As already stated, the fidelity of the GTD has to be known to successfully apply the experimental verification method. To determine the GTD fidelity in a more meaningful way than by the individual measurement errors, we implement the validation framework with a novel reference scene called Normbox and the reference renderer Mitsuba [55]. Assuming that the render error is negligible for our simple reference scene, the difference from the comparison path will comprise the measurement, processing, and comparison errors. A spectral error analysis enables further restriction of the errors' influence. It allows for a spectral localisation of the errors and, consequently, a more substantiated attribution of their sources. In the following, we first present the Normbox and then the implementation of each step of the comparison path.

3.3.1 Normbox

The basis of our validation framework is the reference scene. We introduce the Normbox, a modified Cornell box, as illustrated in Figure 3.2. The Normbox consists of a wooden box, a sample holder, an integrating sphere combined with an opal glass panel for the light source, and a portable spectroradiometer mounted on an aluminium profile to acquire the reference data. Each component is described in more detail below.

Wooden box: The wooden box is a cube made of 15 mm thick multiplex panels and a 19 mm panel for the ceiling. The inner dimensions of the cube are $600 \times 600 \times 600$ mm. One side of the cube is open to allow measurements inside the box. On the ceiling, there is an aperture of dimensions 220×200 mm, where the light of the integrating sphere enters the box. The the box's panels are painted with a matte, gray dispersion paint (RAL 7004) that provides a flat reflectance spectrum with $a^*, b^* < 1$ and $L^* \approx 65$. According to EN ISO 3668:2001 [70], this is sufficient for visual comparison purposes. A black-painted sample holder is placed in the middle of the box. The sample is mounted at the holder's flexible head, enabling its rotation around the horizontal axis.

Area light source: Measuring the emission behaviour of arbitrary light sources is a complex procedure requiring special equipment. A common way to simplify the simulation of the emission behaviour is to approximate it by a homogeneous panel radiator with a Lambertian emission behaviour. To minimise the difference between the assumed light source in the simulation and the used light source




Figure 3.2: The Normbox consists of a wooden box and the spectroradiometer CS-1000A by Konica Minolta. An integrating sphere combined with an opal panel is used as a light source.

in the Normbox, we use the lightbox LE7-4x by Image Engineering ¹ combined with an opal panel. The lightbox is an integrating sphere combined with LED modules producing freely selectable emission spectra, which are kept constant by an internal spectrometer. The homogeneity and diffusivity of the area light source deviate from the Lambertian assumption by 3% and 2.6%, respectively.

Portable spectroradiometer: We use the portable high-precision spectroradiometer CS-1000A by Konica Minolta (Sec. 2.3) to acquire the spectral radiance of the target samples within the Normbox. These spectra constitute the reference data described in the validation framework (Fig. 3.1). As shown in Figure 3.2, an aluminium profile base connects the spectroradiometer with the wooden box. The distance and height between the spectroradiometer and the sample holder are adjustable.

3.3.2 Comparison path

Referring to Section 3.2, the comparison path of the validation framework consists of four steps: GTD acquisition, processing, rendering and comparison. The implementation of these steps with the Normbox and the reference renderer Mit-

¹<https://www.image-engineering.de/products/equipment/illumination-devices/378-le7>

suba is presented in the following.

Step 1: Acquisition of ground truth data

The GTD consists of an accurate description of the Normbox and the reference data. The description comprises the emission behaviour of the area light source in the ceiling, the reflection behaviour of the grey walls and the 24 patches of the x-rite ColorChecker ², as well as the geometries of all involved objects. The reference data includes the spectral radiance of all ColorChecker patches under 13 different angles measured within the Normbox.

Light: As aforementioned, the light source of the Normbox is assumed to be ideally homogeneous and Lambertian. However, the spectral radiance is measured at 25 points to compensate for the light source inhomogeneity. Test renderings show no noticeable differences between the representation with 25 spectral radiance values and one averaged spectrum. Nevertheless, we used in this work the representation with 25 spectra.

Materials: The x-rite ColorChecker covers the sRGB colour gamut and includes six neutral patches. The patches and the grey wall paint have a rough, isotropic, and uniform surface with a noticeable back-scattering. As noted by Rump et al. [71], the ColorChecker patches are far from being ideally diffuse. Hence, we decided to measure the in-plane BRDFs using our gonioreflectometer with the first configuration described in Chapter 2.2.1. Each in-plane BRDF is measured for the incident angles $\theta_i = -80^\circ, -70^\circ, \dots, 0^\circ$ and the reflection angles $\theta_o = -80^\circ, -70^\circ, \dots, 80^\circ$. The incident angle varies only from -80° to 0° since, for isotropic materials, the remaining range is redundant. Note that we did not measure the reflectance behaviour of the black sample holder in the centre of the Normbox since test simulations confirmed that even a black material with a strong specular reflection has a negligible impact on the appearance of the ColorChecker patches. Thus, for simplicity, we assumed the sample holder as diffuse black.

Geometries: The positions and dimensions of the wooden box, the area light source, the sample holder, and the spectroradiometer are measured manually with a yardstick. To adjust the sample angle related to the ground, we used an Apple iPhone 5.

²<https://www.xrite.com/categories/calibration-profiling/colorchecker-classic>

Reference data: The reference data comprises the spectral radiance of all 24 ColorChecker patches under 13 different sample angles measured within the Normbox. The sample angle varies from 30° to 90° in 5° steps, where 0° is parallel and 90° perpendicular to the light source. Hence, the reference data covers direct- and indirect light dominated illumination situations. We acquired the reference data with the spectroradiometer CS-1000A that points towards the centre of the sample holder. In total, we captured 312 spectra in 13 series, where in each series, the sample angle is first adjusted, and then the spectral radiance for all patches is measured. This measurement procedure guarantees the same sample angle for all patches.

Step 2: Processing

The GTD must be processed so that the Mitsuba renderer can import the Normbox’s description data. There are three processing steps necessary. First, the measured geometries are converted into the widely used OBJ file format. Second, the spectral radiance of the area light source is interpolated to the Mitsuba spectral resolution of 200 samples between 360 and 830 nm. And third, an analytical BRDF model is fitted against the measured in-plane BRDFs. The first two steps are self-explanatory, so we will not explain them further. On the contrary, the fitting process is more complex and described below.

BRDF fitting: We approximate the reflection behaviour of the ColorChecker patches and the grey wall with the Cook-Torrance BRDF model [29] along with the GGX normal distribution function and Smiths shadowing-masking function. We refer the reader to Chapter 2.1.4 for the respective definitions. We use the Cook-Torrance model because it is one of the most popular reflection models, and it successfully fits a wide range of measured BRDFs [44, 72].

We observed for some ColorChecker patches a global offset when fitting the Cook-Torrance model. Therefore, the native Cook-Torrance model is slightly modified by subtracting the offset c :

$$f_c(\lambda, \omega_o, \omega_i) = f_r(\lambda, \omega_o, \omega_i) - c. \quad (3.1)$$

For some patches, the offset minimises the fitting error, while for others, it assumes the value zero; hence is not taken into account.

The fitting process is divided into a nonlinear and linear optimisation process, where the latter is a subprocess of the nonlinear optimisation process. The nonlinear optimisation determines the model parameters roughness α and refractive index η , as well as the offset c ; the linear optimisation of the wavelength-

dependent specular albedo k_s and the diffuse albedo k_d using linear regression.

The fitted BRDF f_{fit} and the measured BRDF f_{ref} are the input values for the cost function, which is based on the cost function of Löw et al. [36], and defined as:

$$g(\lambda, \omega_o, \omega_i; p) = \ln(1 + \cos^w \theta_i f_{fit}(\lambda, \omega_i, \omega_o; p)), \quad (3.2)$$

$$\hat{g}(\lambda, \omega_o, \omega_i) = \ln(1 + \cos^w \theta_i f_{ref}(\lambda, \omega_i, \omega_o)), \quad (3.3)$$

$$E(p) = \sum_k \sum_l (g^{kl}(p) - \hat{g}^{kl})^2, \quad (3.4)$$

where k and l represent the incident and reflection directions.

In Equation 3.2 and 3.3 the measured and fitted BRDFs are weighted by the cosine of the incident angle. Consequently, the low diffuse reflection at small incident angles is weighted stronger than the comparatively high Fresnel reflection at grazing angles. The optional exponent $w \geq 1$ further enhances this weighting. For the ColorChecker patches, w was set to values between 1 and 4, and for two patches, to 6. These exponent values improve the fitting of the diffuse component, which is, in our work, more important than the high Fresnel reflection. Furthermore, we observed that the exponent improves the overall fitting error in some cases.

Step 3: Rendering

We use the Mitsuba renderer [55] as the reference renderer. We built it on a 64-bit Windows 10 platform and configured it as a spectral renderer with 200 wavelength samples. We added a BRDF plugin that implements the Cook-Torrance model as defined in Chapter 2.1.4. Furthermore, the bidirectional path tracer and the low discrepancy sampler with 512 samples per pixel are used. With the modified Mitsuba distribution and the processed GTD, we generated spectral images of all ColorChecker patches under 13 angles within the Normbox. Each spectral image has a resolution of 10×10 pixels, where each pixel stores a radiance spectrum. The spectral image is averaged and, henceforth, called predicted data.

Step 4: Comparison

There are different methods to compare the predicted and reference data. In this work, we chose a physical and a perceptual comparison method. The physical method determines the radiance differences between the predicted and the reference data, and the perceptual method provides information about the perceived colour differences.

Physical comparison: The radiance difference is calculated by subtracting the predicted- from the reference radiance spectrum. We observed that the radiance differences scale with the material reflectance and illumination intensities, which hinders a reasonable comparison between different sample angles and colour patches. Therefore, we introduce the normalised spectral error (NSE), where the reference spectrum normalises the spectral difference. This error metric allows for spectral analysis of occurring errors. However, it is not appropriate to determine the total error between two spectra; hence it is difficult to compare the overall difference of all ColorChecker patches with one another. Therefore, the root-mean-square error (RMSE) between the predicted and reference spectrum is additionally calculated. Again, the RMSE is normalised by the reference spectrum to allow for comparability. This error metric is called normalised root mean square error (NRMSE).

Perceptual comparison: A widespread metric quantifying the perceived colour difference is the CIE Delta E metric as reviewed in Chapter 2.1.2. The CIE recommends the latest version of the CIE colour difference metric, the ΔE_{00}^* , for industrial colour evaluation. The ΔE_{00}^* between the predicted and reference spectrum is calculated according to the definition given in the book Color Imaging [13, p.461]. The required tristimulus values are normalised with the reference tristimulus values of the white patch under the sample angle of 30° . This normalisation ensures that all Y values are within 0 and 1. For the transformation from the XYZ- to $L^*a^*b^*$ -space, we use the standard illuminant D65.

3.3.3 Error path

The error path describes the summation of all errors occurring in each step of the comparison path resulting in the total error. According to the four steps of the comparison path, these errors are divided into measurement, processing, rendering, and comparison error. Each group is subdivided into various sources of errors as shown in Figure 3.3.

Step 1: Measurement error

The measurement error is divided into four different errors introduced during the GTD acquisition. The first two regard the geometric measurements of the Normbox. As mentioned in Section 3.3.2 the geometric dimensions are measured with a yardstick with a measurement error of approximately ± 1 mm. The sample angle is adjusted using an Apple iPhone 5 with an accuracy of roughly $\pm 1^\circ$. The third measurement error regards the in-plane BRDF measurement. This error is

Measurement Error	Processing Error	Rendering Error	Comparison Error
<ul style="list-style-type: none"> - Geometries Normbox - Angle of target sample - In-plane BRDF 	<ul style="list-style-type: none"> - Conversion to obj format - Fitting BRDF model 	<ul style="list-style-type: none"> - Assumption: <i>For Mitsuba negligible</i> 	<ul style="list-style-type: none"> - Physical comparison <i>No Error</i> - Perceptual comparison: <i>Non uniformity of dE2000 metric</i>
<ul style="list-style-type: none"> - Light source emission - Reference data 	<ul style="list-style-type: none"> - Repeatability - Interpolate spectrum - Deviations from assumed light source 		

Figure 3.3: Error sources that occurs during the implementation of the comparison path of the validation framework

defined by the measurement error of our gonioreflectometer, which is on average 0.7 % as described in Chapter 2.2.3. Finally, the fourth error source is the spectroradiometer CS1000A, which measures the reference data and the light source emission with an accuracy of 2 %.

Step 2: Processing error

The conversion of the measured geometries into a digital format, such as the popular obj format, introduces negligible errors due to the geometrical simplicity of the Normbox. On the contrary, fitting the BRDF model and approximating the light source emission lead to considerably higher errors. The averaged fitting error varies between 5 % and 10 % depending on the patch. For patch 18, we observed a conspicuous error of 18 %. Note that the fitting error is smaller for the diffuse reflection than for the specular reflection due to the cosine weighted cost function (Sec. 3.3.2).

The processing error of the light source emission is divided into the errors: repeatability, spectrum interpolation, and light source approximation. The repeatability error is <1 %, and the spectrum interpolation error is negligible due to the densely sampled spectrum with 200 samples. The Lambertian assumption leads to an error of 3 % and 2.6 % for the homogeneity and diffusivity, respectively. The inhomogeneity is compensated by measuring the emission spectra at 25 locations. The intensity of the light source is kept constant during all measurements. Notwithstanding, we registered a maximal ΔE_{00}^* of 0.28.

Step 3 & 4: Rendering error & Comparison error

Finally, we consider the rendering and comparison errors. The rendering error is undoubtedly unknown, but we assume that it is negligible for Mitsuba due to the simplicity of the Normbox. The comparison error is different for the physical and perceptual comparison methods. The physical comparison is error-free since we compare the raw radiance spectra. Conversely, the perceptual comparison is

error-prone because the raw data is processed. The ΔE_{00}^* color-difference formula is based on the perceptually non-uniform $L^* a^* b^*$ -space, as explained in Chapter 2.1.2. The formula compensates for some of these non-uniformities but introduces new ones on the flip side. Hence, there is still the problem that the same ΔE_{00}^* can be perceived differently depending on location in the $L^* a^* b^*$ -space.

Results of error path

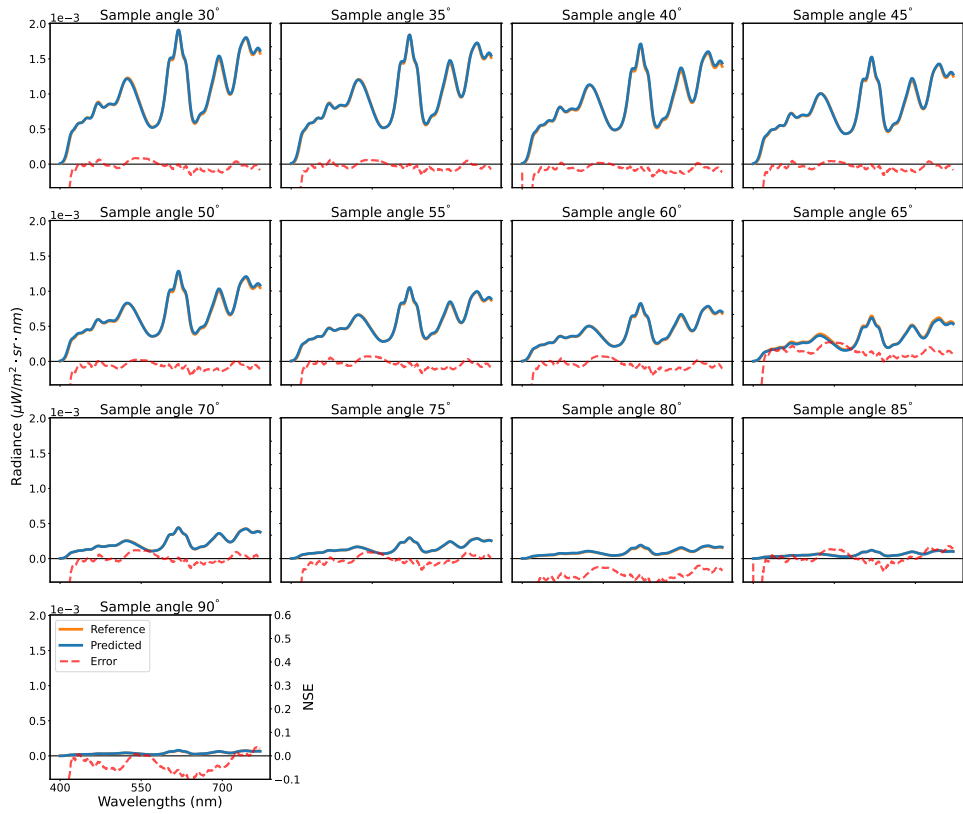
The implementation of the error path demonstrates that, despite the reference scene’s simplicity, many errors contribute to the total error. The quantification of these errors supports our assumption that the material representation, more precisely, the fitting of the in-plane BRDF, is the most critical error. Note that the fitting error even increases when computing global illumination due to multiple light reflections and the accompanying error summation. The remaining errors are much less critical, whereby the light source constitutes, apart from the fitting error, the most significant error source due to the deviations from the Lambertian assumption.

3.4 Fidelity of ground truth data

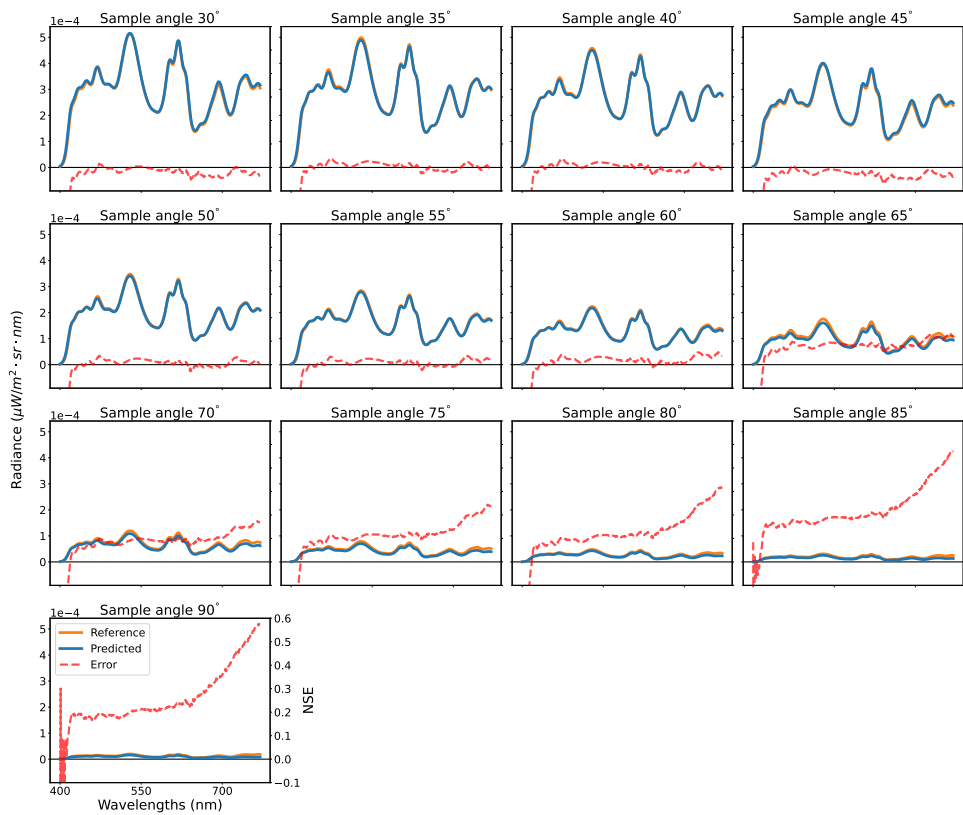
So far, we considered all error sources individually without knowing how strong they affect the GTD’s fidelity. Estimating the error by error propagation is hard, or even impossible, due to the complexity of the render equation. However, knowing the GTD’s fidelity is paramount for the experimental verification method. In the following, we first quantify and analyse the physical and perceived colour differences between the previously acquired reference and predicted data using the error metrics introduced in Section 3.3.2. Then, we compare our results to previous works by Schregle and Wienold [60] and Bärz et al. [64].

Results of physical comparison

Figure 3.4 illustrates the reference and predicted spectral radiance and the respective NSE of Patches 2 and 23. Each subplot contains the result of a specific sample angle. We selected these two patches because they represent patches with a low (Patch 2) and large (Patch 23) NSE. The figure illustrates that for both patches, the reference and predicted spectra match well for small sample angles up to 60° , also confirmed by the low NSE. From a sample angle of 65° the NSE constantly increases, particularly strong for Patch 23. When analysing the wavelength dependency of the NSE, we can observe at Patch 23 a notable increasing trend for large wavelengths.



(a) Patch 2



(b) Patch 23


Figure 3.4: Comparison of the reference and predicted spectral radiance of Patches 2 (a) and 23 (b) and the respective NSE.

Based on these observations, we make the following three assumptions. First, the GTD has a high fidelity due to the consistently small NSE at Patch 2. Second, the considerably higher NSE at patch 23 is primarily caused by the fitting error, hence by the used reflection model. And third, for large sample angles, the proportion of indirect light increases, which leads to larger errors due to multiple reflections and large fitting errors for the specular reflection. In the following, we will further investigate these assumptions by analysing the results of all ColorChecker patches.

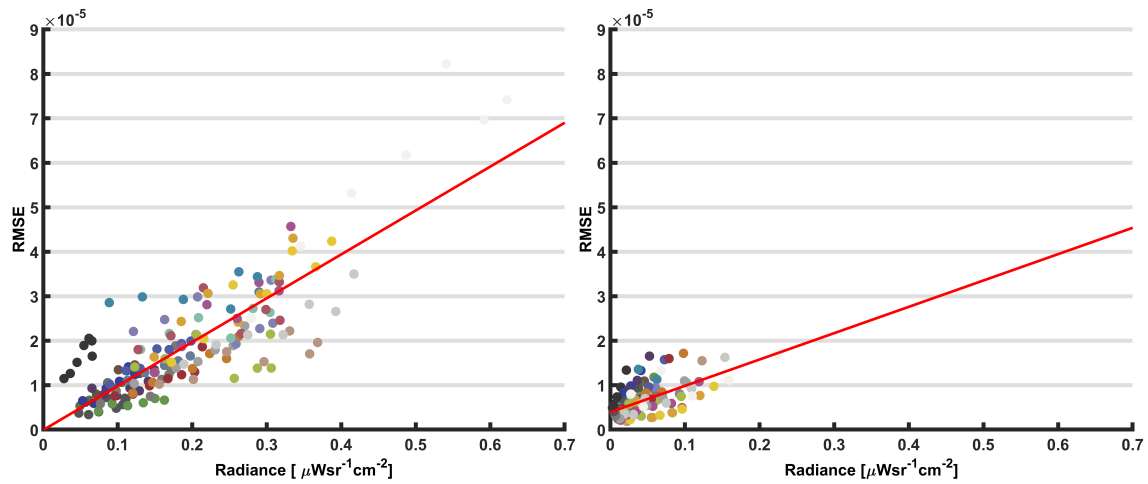
Error quantification: We computed the RMSE between the reference and predicted spectra as described in Section 3.3.2. Figure 3.5 (a) depicts the RMSE for all patches and sample angles. It can be remarked that there is a significant RMSE variation between different patches and sample angles. In particular, patch 19 has a prominent high RMSE. When comparing the grey patches (Patches 19-23, excluding Patch 24), there is a correlation between the patch’s radiance and RMSE. This correlation can also be observed when looking at each patch individually. The RMSE decreases for increasing sample angles, hence decreasing radiance.

The observed correlation between the radiance and RMSE is further examined in the scatter plots shown in Figures 3.5 (b) and (c). Plot (b) shows the data points of all patches for small sample angles from 30° to 60° and plot (c) for the remaining large sample angles from 65° to 90° . Plots (b) and (c) illustrate a strong and moderate positive linear relationship between the radiance and RMSE, with correlation coefficients of 0.84 and 0.62. In brief, there is a positive linear correlation between the radiance and RMSE, particularly for small sample angles. This correlation makes sense since, for all patches, the illumination is the same, and the relative fitting error is constant, except for some outliers. Thus the patches with high reflectivity lead to larger absolute errors (i.e. higher RMSE) than patches with low reflectivity.

We calculate the NRMSE to remove the scaling factor, thus, comparing the results of different measurements more meaningfully. As shown in Figure 3.6 (a), the NRMSE of all patches, except patch 24, is almost on the same level for small sample angles. But at larger sample angles, there is an increasing trend of the NRMSE, which even varies between different patches. We again plotted the data points for small and large sample angles in separate scatter plots as illustrated in Figure 3.6 (b) and (c). In plot (b), nearly all data points have an NRMSE between 0.02 and 0.05, except for patch 24 and two data points of patch 18. The scatter plot shows that for small sample angles, there is no correlation between radiance and NRMSE, per our observations from the plot (a). For large sample




(a) RMSE between the reference and predicted data of all patches and sample angles (bar color).



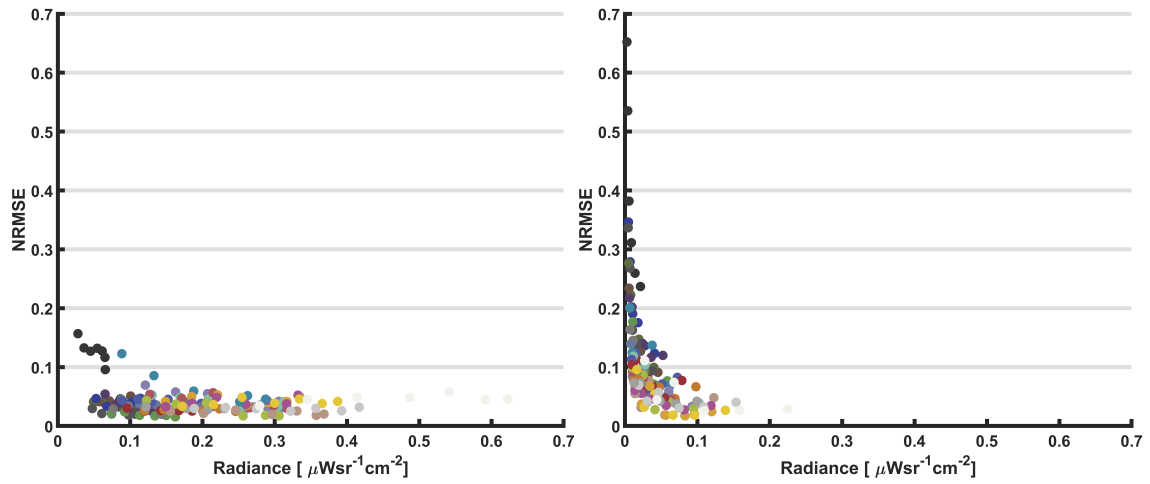
(b) Correlation between radiance and RMSE of small sample angles (30° - 60°)

(c) and large sample angles (65° - 90°).

Figure 3.5: Investigation of the correlation between the sample radiance and the RMSE. A clear positive linear relationship with a correlation coefficient of 0.84 and a moderate linear relationship with a correlation coefficient of 0.62 can be observed for small and large sample angles, respectively.



(a) NRMSE between the reference and predicted data of all patches and sample angles (bar color).



(b) Correlation between radiance and NRMSE of small sample angles (30° - 60°)

(c) and large sample angles (65° - 90°).

Figure 3.6: Investigation of the correlation between the sample radiance and the NRMSE. There is no correlation for small sample angles; contrary, an exponential increase of NRMSE for decreasing radiance can be observed for large sample angles.

angles, the NRMSE increases exponentially for decreasing radiance, as depicted in the plot (c). Furthermore, the NRMSE tends to be higher for bluish and greenish patches than for reddish and yellowish patches. Summarised, the NRMSE for small sample angles is nearly constant while exponentially dropping with increasing radiance for large sample angles. We further analyse these observations in the following section.

Error analysis: As stated in Section 3.3.2, the 13 sample angles describe a transition from direct- to indirect-light-dominated illumination. Moreover, we know that the light composition for sample angles up to 60° is roughly the same, while

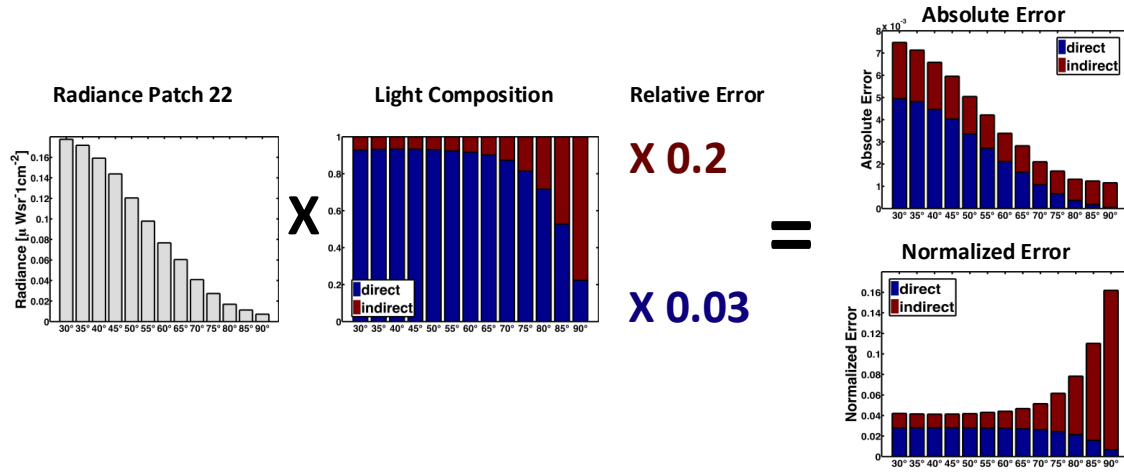


Figure 3.7: Error model: The absolute error is simulated by multiplying the radiance with the light composition and the two errors, err_d (blue) and err_i (red). err_d is caused by the fitting error of the patch; err_i by the fitting error of the grey wall.

after 65°, the proportion of indirect light increases exponentially. The increasing NRMSE with increasing sample angles, as illustrated in Figure 3.6 (a), correlates with the increase of indirect light. The main difference between indirect and direct light is the additional reflections at the grey walls; hence the additional fitting error is a reasonable explanation for the increasing NRMSE. However, the NRMSE also varies from patch to patch, as illustrated in Figure 3.6a. Since the indirect illumination is for all patches the same, it is very likely that, besides the grey wall’s fitting error, the ColorChecker patches’ fitting error greatly impacts the overall error. We suspect that the material representation, more precisely, the fitting of the BRDF, is the dominant error. Moreover, we assume that the overall difference between reference and predicted data can be approximated by two errors; the first error is related to direct light (err_d), and the second error is to indirect light (err_i). The fitting errors define both errors.

In our setup, the sample angles avoid that direct light is simultaneously incoming and measured from grazing angles. Thus, the measured direct light is not affected by the observed high fitting errors for the specular reflection. On the contrary, indirect light comprises all scattering events, single and multiple reflections, occurring within the Normbox. Hence, it is affected by all fitting errors of the grey wall, including the high fitting errors at grazing angles. Furthermore, the indirect light leads to grazing incident angles, where fitting errors of the ColorChecker patches are generally higher than under small incident angles. Therefore, we know that err_i is much larger than err_d .

Based on these insights, we simulate the error of Patch 22 as shown in Figure 3.7. In the simulation, the radiance of the reference data is weighted by the

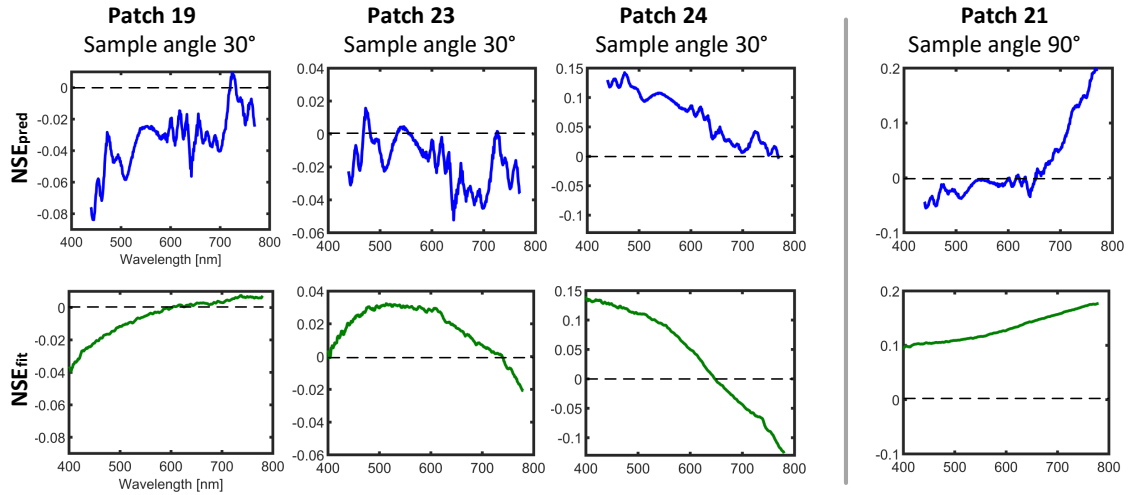


Figure 3.8: Top row: NSE_{pred} for patches 19, 23 and 24 at sample angle 30° (left), and for patch 21 at sample angle 90° (right). Bottom row: respective NSE_{fit} .

light composition to determine the proportion resulting from direct and indirect light. Both proportions are multiplied with the respective error (err_d , err_i) to compute the absolute error. err_d is defined by the averaged fitting error of Patch 22 at the respective incident and reflection angles of the direct light. And err_i is estimated by means of the averaged fitting errors of the grey wall and Patch 22. The predicted absolute and normalised errors are illustrated on the right side of Figure 3.7. Both resemble the RMSE and NRMSE in Figure 3.5a and 3.6a, which supports our assumption that the fitting error is the dominant error.

To further verify this assumption, we compared the NSE between the reference and predicted data (NSE_{pred}) with the NSE between the measured and fitted BRDF (NSE_{fit}). Figure 3.8 depicts in the top row the NSE_{pred} for Patches 19, 23, and 24 at the sample angle of 30° , and in the bottom row, the respective NSE_{fit} at incident and reflection angle of 30° and -60° . The figure demonstrates that the NSE matches well for Patches 19 and 24 and moderately for Patch 23.

On the right side, Figure 3.8 illustrates the NSE_{pred} of Patch 21 at the sample angle of 90° , and the NSE_{fit} of the grey wall at grazing incident and reflection angles of -70° and 80° . We compare these errors because we assume that the grey wall's fitting errors at grazing angles dominate when illuminating with indirect light. The NSE_{pred} and NSE_{fit} of Patch 21 show a similar increasing trend towards large wavelengths. For both small and large sample angles, the spectral distribution of the NSE_{pred} is predominantly characterised by the NSE_{fit} , which again supports our fitting error hypothesis.

The error quantification and analysis confirm all three assumptions made at the beginning of this section. First, the low NRMSE for several patches and the restriction on the processing error verify the low measuring errors, hence the high

fidelity of the GTD. Second, the error simulation assuming two separate errors related to direct and indirect light demonstrates that the observed exponential increase of the NRMSE for large sample angles stems from the increasing proportion of indirect light. And finally, both the error simulation and the spectral error analysis confirm the fitting as the dominant error source.

Fitting error analysis: We analyse the fitting error in more detail to better understand why the Cook-Torrance model cannot adequately represent the measured in-plane BRDF of the ColorChecker patches and the grey wall. We continue analysing the fitting error of the grey wall since it has a uniform reflectance spectrum and a similar reflection behaviour as the ColorChecker patches. Figure 3.9 illustrates the measured and fitted BRDF for three different scattering phenomena. The three plots depict back-scattering on the left, diffuse reflection in the middle, and forward scattering on the right. The incident angle is in all scenarios -80° and the reflection angles are -80° , 0° and 70° .

The diffuse reflection is fitted well by the diffuse term of the Cook-Torrance model. As expected, the back-scattering peak is not well represented since it is usually caused by multiple reflections, which the model approximates with the Lambertian diffuse term; hence the angle dependency of the back-scattering is not considered. For the fitting errors in the forward scattering, we have no reasonable explanation.

When analysing the spectral distribution, we can observe wavelength shifts that are new to computer graphics. The diffuse BRDF spectrum is, as expected for grey wall paint, uniform. However, the forward and back-scattering spectra are tilted towards long and short wavelengths. These wavelength shifts don't originate from measurement errors, as proven in Chapter 2.2.3. So far, in computer graphics, models based on geometric optics assume that the surface reflection is entirely described by the Fresnel equations and, consequently, for di-




Figure 3.9: Measured and fitted BRDF of the gray wall for three different reflection angles. From left to right, the plots show BRDF spectra which are dominated by back-scattering, diffuse reflection, and specular reflection.

electrics, almost wavelength independent. Figure 3.9 demonstrates that this assumption is incorrect and leads to perceptual deviations from reality. The Cook-Torrance model fails to fit such wavelength shifts adequately, particularly in the back-scattering. In the forward scattering, the model somehow fits the linear increasing spectrum using the physically not plausible specular albedo.

Results of perceptual comparison

The physical comparison provides valuable information about the GTD’s fidelity and the influence of the individual error sources. But the physical comparison does not give any information about the perceived colour difference between the reference and predicted data. Nevertheless, this difference is important for colour-critical applications of rendered images such as VP. In Figure 3.10 we present the color differences E_{00}^* for all patches and sample angles. In contrast to the RMSE and NRMSE in Figures 3.5a and 3.6a, the perceived color differences do not show any trends or patterns. For most patches, the ΔE_{00}^* are below the discernible threshold for a next-to-each-other comparison. Only patches 10, 13, 19, and 24 have values slightly above two.

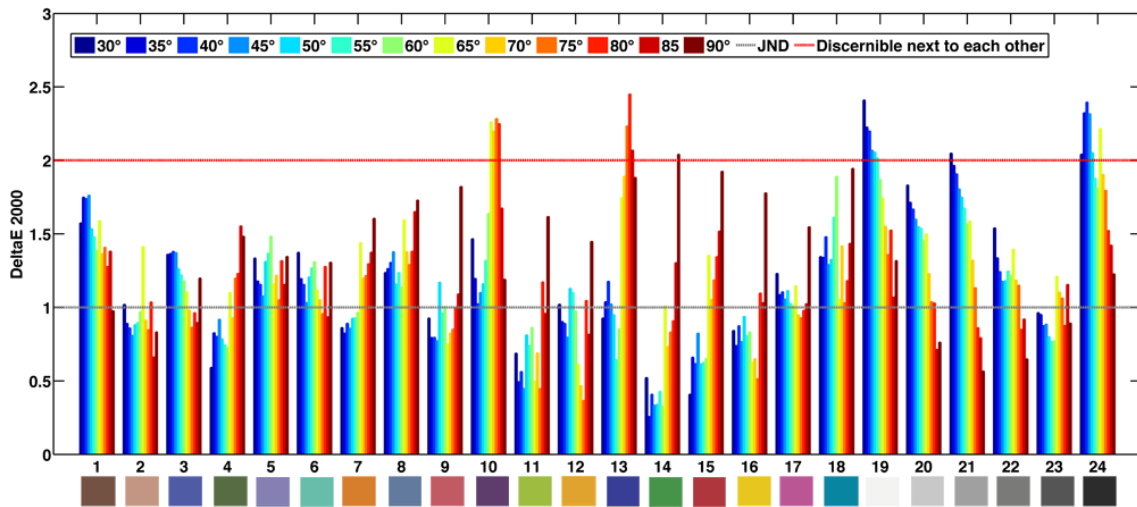


Figure 3.10: Color difference (DeltaE 2000) between reference and predicted data of all patches and sample angles.

Ideally, all data points would have a ΔE_{00}^* below the JND of one; consequently, there would be no noticeable difference between reference and predicted data in a side-by-side comparison. However, in our implementation, this is often not the case. Although we used a simple reference scene and the reference renderer Mitsuba, there are still perceivable differences between reference and predicted data. As we found out in the physical comparison, these differences mainly originate from the GTD’s processing, more precisely, from modelling the material

reflectance. Consequently, improving the material modelling is of great importance for improving the fidelity of VP.

Comparison to previous work

We compare our results with those of Schregle and Wienold [60] and Bärz et al. [64]. Both are the most recently published works implementing the experimental verification method and provide outstanding small differences between reference and predicted data.


Schregle and Wienold conducted a physical comparison between reference and predicted illuminance values. They observed a mean deviation of 2% in their case study for diffuse patch reflections, where they investigated the light reflected from a diffuse grey molleton patch. With our approach, most patches have deviations below 2%, and several patches have significantly lower deviations, as illustrated in Figure 3.11a. Considering patch 20, which has a similar reflection behaviour to the grey molleton, our approach leads to a better result.

Bärz et al. conducted a perceptual comparison. They computed the colour difference ΔE_{76}^* between the reference and predicted data of all ColorChecker patches for one specific sample angle. We compared our results at the sample angle of 45° with their results, as shown in Figure 3.11b. The figure demonstrates that the colour differences of our approach are sometimes higher than the ones of Bärz et al.. However, our results are more stable and do not contain extreme outliers, as observed for Patches 13 and 15.


Our GTD has a similar or even better fidelity than previous works. Moreover, we provide for the first time spectral GTD enabling a more detailed analysis of arising deviations between reference and predicted data.

3.5 Conclusion

The light simulation has a great influence on the appearance of virtual prototypes. To improve VP, it is inevitable to validate the light simulation to identify its bottlenecks. We introduced a validation framework for the light simulation of render software. The framework implements the experimental verification method; hence ground truth data comprising the accurate description of a reference scene and reference data are required. We developed a novel reference scene (Normbox) that provides controlled conditions. The key feature of the Normbox is the use of an integrating sphere which is controlled by an internal spectrometer. Combined with an opal panel, it is a stable and nearly ideal homogeneous and diffuse area light source. Furthermore, the Normbox supports the measure-



(a) Comparison of our approach with the one of Schregle and Wienold (red dashed line): Relative illuminance differences between reference and predicted data for diffuse reflection.



(b) Comparison of the color differences (DeltaE 76) between reference and predicted data with the approach of Bärz et al..

Figure 3.11: Comparison to related work.

ment of the spectral radiance of samples placed inside the box. We acquired a set of precisely and spectrally resolved GTD. The GTD consists of the Normbox description, including the in-plane BRDFs of 24 ColorChecker patches and reference data comprising 312 radiance spectra of all patches under 13 angles. The reference data covers rough, dielectric materials with different reflectance spectra and various illumination situations, from direct- to indirect-light-dominated scenarios. We implemented the framework using the Cook-Torrance BRDF with GGX distribution as a fitting model and Mitsuba as a reference renderer. We compared the resulting predicted and reference data with a physical and perceptual method. The physical comparison confirmed that the fidelity of the GTD is high and that the fitting of the measured BRDFs is the dominant error. Furthermore, we observed unusual linear wavelengths in the measured BRDFs shifts towards short and long wavelengths in the back- and forward scattering. These colour shifts are new to computer graphics, consequently, not considered by the used Cook-Torrance model. The perceptual comparison demonstrated that these errors lead to perceptual colour differences between predicted and reference data. Nevertheless, our comparison leads to superior results when compared to previous works. Physically based rendering can achieve photorealism but fails to be predictive. Current BRDF models cannot reproduce all perceptual critical scattering phenomena; thus, more research in this area is needed to improve the fidelity of VP. Our GTD is a good starting point to further investigate the influence of the employed approximations on the appearance of real materials. In this work we used the Normbox, which provide a very simple geometric and light situation, so that we can assume the Mitsuba rendering error as negligible. To validate sophisticated rendering techniques a more complex reference scene is required where the state of the art integration and sampling algorithms struggles.

4. Light interaction with rough surfaces

In the previous chapter, we demonstrated that the simulation of light-matter interaction is a major challenge in computer graphics. Often, it is described by an analytical formulation, usually by physically based BRDFs. These BRDFs are designed based on real-world observations or physical measurements, and BRDF databases, such as the well-known MERL database [17], are of important value. Nevertheless, these databases are not free of drawbacks making a sound analysis and modelling of light-matter interaction hard or, in some situations, even impossible. For the MERL database, for example, the light spectrum is represented by RGB triplets, complicating modelling wavelength-dependent scattering phenomena.

Current BRDF models based on geometric optics lead to photorealistic images but often fail to be predictive, especially for rough surfaces, as we demonstrated in the previous chapter with the ColorChecker patches. Such surfaces contain several scales of roughness where many distinct scattering phenomena occur. These are particularly challenging to model and cannot be described consistently by geometric optics alone.

Recently, more sophisticated BRDF models based on wave-optics were presented as reviewed in Chapter 2.1.4. They are capable of modelling wave-optics phenomena and lead to a more realistic appearance. However, they are still limited to some approximations, e.g. ignoring multiple reflections. And, to our knowledge, for rough surfaces, there is no sound validation with measured data. Both limitations make it difficult to classify their agreement with reality which, in turn, hinders further improvement of the models. Finally, the computational cost of these models is another disadvantage making them unsuitable for real-time applications.

A deeper understanding of wave-optics phenomena is essential to overcome the limitations and further improve the simulation of light interaction with rough surfaces from visual and performance perspectives. Therefore, in this chapter, we experimentally investigate this interaction by generating aluminium samples

with varying roughness, whose optical properties are captured with a multi-modal measurement approach, including: surface topography from a confocal microscope, measured spectral in-plane BRDF, macro photos under varying illumination angles, and, finally, scattering simulation with a virtual goniometer. Such a rich dataset enables a deeper analysis of light interaction with rough surfaces and, consequently, to validate or even improve existing models.

After a thorough analysis of the acquired data, we found that diffraction takes place in all cases and actually dominates the appearance of the aluminium samples. We show that wavelength shifts, leading to a reddish and bluish appearance, can be related to forward and back-scattering. A linear function can surprisingly describe these wavelength shifts. These contributions are strengthened by the fact that data from all different measurement modalities and wave and geometric optics simulations corroborate our claims. Moreover, the detailed analysis and the proposal of a simple model pave the way to more complete BRDF models able to simulate diffraction effects without impacting the computational costs.

4.1 Related work

To understand and analytically describe the light interaction with materials, measured BRDFs are paramount. Important databases and their limitations are presented in the following. Further, the limitations of BRDF models based on wave-optics are summarised.

Databases

The MERL database [17] provides densely sampled BRDFs of many materials and has been widely used to develop and validate reflectance models. The BRDFs are acquired with a camera that drastically shortens acquisition time but, on the flip side, only provides RGB values and introduces lens aberrations [18]. The UTIA [73, 74] and the UBO2014 database [75] provide BTFs of many textured materials, but both are also limited to RGB measurements. Furthermore, all scanned materials are inhomogeneous, making a thorough analysis of the light-matter interaction challenging.

Dupuy and Jakob [24] recently published a spectral database where BRDFs are measured with a dynamic sampling strategy that significantly decreases measurement time. It is very beneficial for spectral BRDF validation. Nevertheless, additional information on the measured materials, such as surface topography, is necessary to analyse the light-matter interaction.

In the optics community, Donnell and Mendez [76] and Schröder et al. [77]

acquired the surface topography and scattering of defined metal samples with varying roughness. Both measured the scattered data at three wavelengths; in both cases, two wavelengths are outside the visible light spectrum.

None of the described databases provides, at the same time, dense spectral BRDF data and information about the surface topography. In this work, we created a dataset that combines many of these essential aspects to analyse the relation between light-matter interaction and surface roughness.

Limitations of BRDF models

Computer graphics commonly describes the light-matter interaction with BRDF models based on geometric optics, which heavily simplify the computations, consequently allowing for usage in real-time applications. Albeit their advantages, deviations can still be noted when comparing or fitting them against measured data. In particular, they often do not adequately represent wavelength-dependent scattering phenomena, such as the wavelength shifts observed in Chapter 3. This issue has led to physical, but not plausible, extensions to improve fitting results, e.g. the use of a specular albedo or two lobes [18, 29, 36, 44].

One reasonable explanation for the deviations between the model and measured data is diffraction, as stated by Holzschuch and Pacanowski [46]. Yan et al. [3] also noticed visual prominent diffraction phenomena on rough surfaces even when illuminated with partially coherent light. The wave-optics models presented in Chapter 2.1.4 enable the simulation of wave-optics phenomena, but they are computationally expensive and not systematically validated with measured data. Even in the optics community, only a few works validate the underlying wave-optics approaches, and for the popular GHS theory, usually used for rough surfaces, Schröder et al. [77] noticed deviations between measured and simulated data. The authors assume that these deviations result from measurement errors, but in their work it is not conclusively clarified.

The work by Yan et al. [3] is particularly important for our analysis of wave-optics phenomena on rough surfaces. As reviewed in Chapter 2.1.4, they derived a generalised BRDF representation including three scalar diffraction theories that allow the simulation of full diffraction effects of arbitrary micron-scale height-field geometries. The authors provide interesting renderings of a brushed metal patch with a clear reddish appearance in the forward scattering. Even though not discussed by the authors, this strongly reminds us of the wavelength shifts observed in the previous chapter at the ColorChecker patches.

4.2 Experimental procedures

The scattering of light on rough surfaces is a complex and high-dimensional problem. The surface's roughness plays a decisive role in the interaction of light and, consequently, the material's appearance. To further investigate this relation, we generated aluminium samples with varying roughness, thus ruling out differences in material properties in our analysis. We used four measurement techniques to conduct a sound analysis: surface measurements, macro photos, spectral BRDF measurements, and light scattering simulations. The aluminium samples and measurement methods are further described below.

Aluminium samples

We generated 8 round aluminium samples since it is a cheap material and easy to process. It provides a high and almost uniform reflection over the light spectrum, and for an ideally smooth surface, the Fresnel equations describe the light interaction well. The required spectral refractive index and extinction coefficient of aluminium are available in the literature [78].

We use the aluminium alloy AlMg3, which is more corrosion-resistant than pure aluminum [79]. While the reflectance of AlMg3 is slightly lower than that of pure aluminium, the spectral shape remains similar. We assume these deviations are small compared to the investigated scattering phenomena and do not negatively influence our analysis.

All samples are 10 mm thick with a diameter of 50 mm and were first polished to a mirror-like finish. Then, all samples except one were sandblasted under varying conditions. We modified the following conditions for each sample: pressure, distance, and abrasives diameter. The chosen parameter values are specified in Table 4.1.

As expected, increasing the pressure and decreasing the distance leads to a rougher surface. Note that using smaller beads as abrasives does not lead to a smoother surface, probably because the smaller beads are accelerated faster than the larger beads at the same pressure.

Surface measurements

As aforementioned, the surface finish impacts the scattering of light and, thus, the material's appearance. It is of great importance when dealing with wave-optics phenomena. Therefore, we measured the surface topographies of all samples with the white light confocal microscope Nanofocus μ surf Generation C. For the device specifications, we refer the reader to Chapter 2.3.

sample No.	pressure (bar)	distance (mm)	abrasives (μm)	roughness RMS (μm)
1	6.5	100	70-100	2.39
2	4	100	70-100	1.76
3	3	100	70-100	1.20
4	2	100	70-100	0.98
5	1	100	70-100	0.52
6	-	100	-	0.10
7	1	200	35	0.66
8	1	200	70-100	0.43

Table 4.1: Sandblasting conditions to generate the eight aluminium samples with varying roughness

Figure 4.1 shows the surface topographies of the eight aluminium samples. For each one, an area of $1280 \times 1280 \mu\text{m}$ was scanned with a vertical resolution of 20 nm for Samples 5, 6 and 8, and of 40 nm for the remaining. The variation in resolution is due to the device’s limitations when scanning rougher surfaces. The figure illustrates that the samples’ topographies consist of randomly distributed spherical holes. Except for Sample 8, where the base surface is still visible, all others show an isotropic random surface. Sample 6 was not sandblasted but still contains visible scratches, influencing the scattering behavior [80].

Macro photos

We acquired macro photos of the aluminium samples to capture the spatial diffraction pattern on the aluminium samples. We used the simple goniometer described in Chapter 2.3 with a fix camera position of -45° . The illumination angle varies from -70° to 70° in 5° steps, where the illumination angles -50° and -45° are left out, due to self-shadowing.

BRDF measurements

The macro photos provide spatial information, but the spectral domain is limited to RGB values. Since wave-optics phenomena are strongly wavelength-dependent, spectral BRDFs are necessary. We measured the in-plane BRDF of all samples with our custom-built gonireflectometer along with the second configuration as detailed in Chapter 2.2. The BRDF is sampled with a resolution of 10° in the incident domain (θ_i) and a dynamic resolution in the reflection domain (θ_o). We discarded Sample 7 due to its inhomogeneous surface.

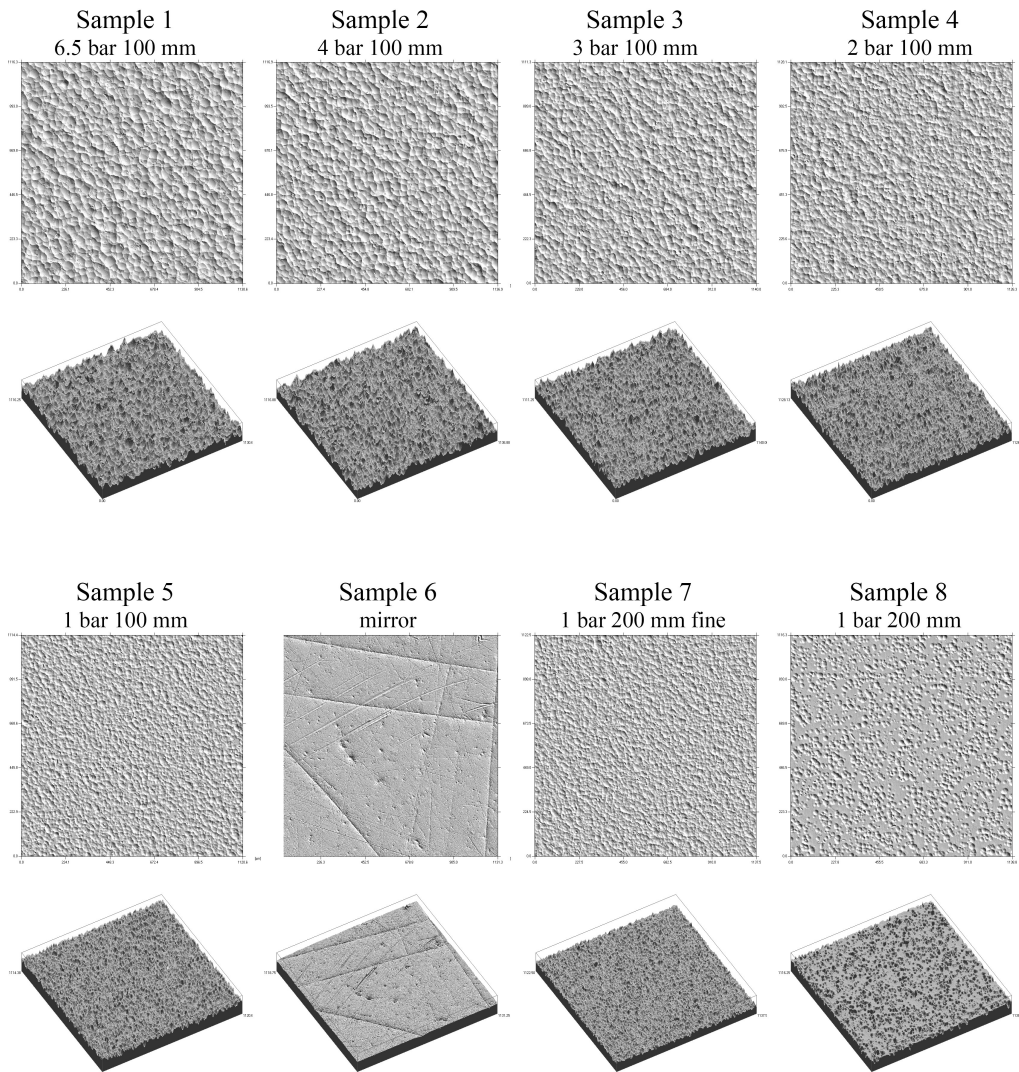


Figure 4.1: Top and 3D views of the samples' surface topographies measured with a confocal microscope. Except for Sample 6, the surfaces consist of randomly distributed holes caused by the glass beads of the sandblaster.

Scattering simulations

Another essential piece of information is the scattering paths, which is particularly important when dealing with rough surfaces where multiple scattering occurs. We used the virtual gonioreflectometer described in Chapter 2.3 to gather information on the energy coming from rays that reach the sensor after a different number of reflections, taking Fresnel into account. Even though the simulation only considers geometric optics, it provides valuable insight into the trends between first and multiple reflections for forward and back-scattering.

4.3 Observations

This section describes the observed light interaction with the rough aluminium surfaces. The overall scatter characteristics of all samples are compared, the visual appearance is evaluated based on the macro photos, and the angle and wavelength dependency is investigated using the in-plane BRDF.

Scatter characteristic

As intended, the aluminium samples scatter light in different ways. Figure 4.2 shows the reflectance distributions at 600 nm of Samples 1, 2, 4, and 6 for various incident angles. The BRDF of the smoothest sample (Sample 6) has a clear peak in the specular reflection and falls quickly to almost zero, resulting in a narrow specular lobe with a high dynamic range of five decades. In contrast, the roughest sample (Sample 1) scatters light in a wide angle range, and the off-specular peak for large incident angles, as reported by Torrance and Sparrow [28], can be observed. Both samples have back-scattering at grazing incident angles, which is stronger at the rough sample. The reflectance distributions of the remaining samples are within these two extremes.




Figure 4.2: Reflectance distribution of Samples 1, 2, 4, and 6 at wavelength 600 nm. Sample 6, with a mirror finish, has a narrow specular peak. Instead, the rougher samples scatter more broadly and present an off-specular peak for large incident angles.

Visual appearance

The macro photos give a first impression of light scattering on rough surfaces. Different from spectral measurements, they provide spatial resolution, which is essential for analysing spatially varying phenomena, such as diffraction.

Figure 4.3 provides macro photos for each sample, where in each case, on the top-left and top-right are, respectively, the macro photos of forward and back-scattering, and on the right side, the corresponding detailed views as well as the average colour of the insets. The figure shows three interesting scattering effects. First, for all samples, except Sample 6, a speckle-like colour pattern can be observed, particularly in the forward scattering. Second, rough samples appear reddish and greyish-bluish in the forward- and back-scattering, which is in line with the observations made from the ColorChecker patches in Chapter 3.4. Third, with increasing roughness, the appearance in the forward scattering change from greyish-bluish to reddish, and in the back-scattering, the other way around.

The observed colour pattern resembles speckle patterns typically observed when rough surfaces are illuminated with coherent light. The speckle effect results from the interference of many waves with varying phases and amplitudes but with the same frequency. A similar effect can also be observed when rough surfaces are illuminated with partially coherent light [3, 81], e.g. natural light. We conclude that these patterns, as well as the on average reddish and greyish-bluish appearance, originate from diffraction.




Figure 4.3: Eight aluminium samples with surfaces varying from smooth (top-left) to rough (bottom-right). For each sample, we show macro photos of forward and back-scattering on the left and a detailed view, as well as the average colour on the right. For all samples, colourful diffraction phenomena can be observed, which is particularly strong for rough surfaces where the appearance strongly deviates from the expected greyish of aluminium.

Angle and wavelength dependency

To further understand the observed diffraction phenomena and to analyse their angle and wavelength dependency, we evaluate the acquired in-plane BRDFs of the roughest sample (Sample 1). In this case, the diffraction phenomena are more prominent, as shown in Figure 4.3, and the low dynamic range simplifies the light scattering analysis. Since the analysis of the in-plane BRDF with almost 300 spectra is complicated, we introduced a simplified RGB representation as illustrated in Figure 4.4. On the x- and the y-axis, the reflection (θ_o) and incident angles (θ_i) are plotted, and each square represents an individual measurement of the in-plane BRDF. To generate such an RGB image, we virtually illuminate the spectra by multiplying them with the standard illuminant D65 spectrum and the incident angle's cosine. The resulting spectra are converted into the sRGB colorspace (Sec. 2.1.2). Hence, a uniform BRDF spectrum results in an RGB value of (1,1,1) perceived as white when adapted to D65.

In agreement with our previous observations, Figure 4.4 (left) shows that the measured BRDF of Sample 1 appears reddish in the forward scattering and greyish-bluish in the back-scattering. A new insight is a continuous transition between the reddish and greyish-bluish appearance.

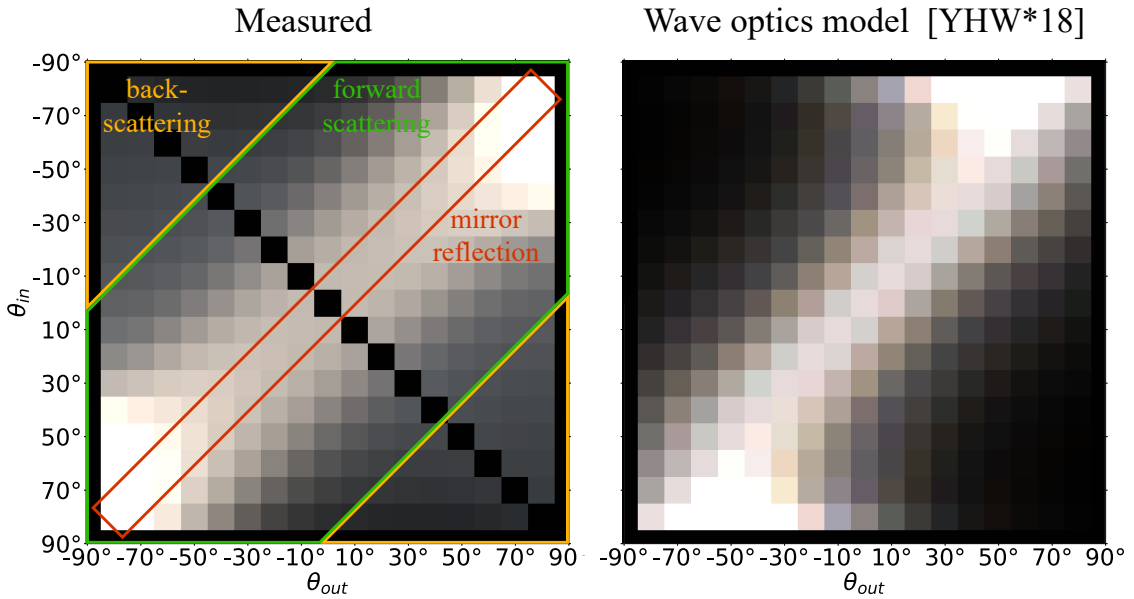


Figure 4.4: In-plane BRDF of Sample 1 (left) depicted in an RGB image and simulated results using the wave-optics model from Yan et al. [3]. x and y-axis are, respectively, reflected (θ_o) and incident (θ_i) angles. The reddish tone in the mirror reflection (centre of the forward scattering) changes continuously into a greyish-bluish tone in the back-scattering region. For the simulated data, the back-scattering effect is not contemplated.

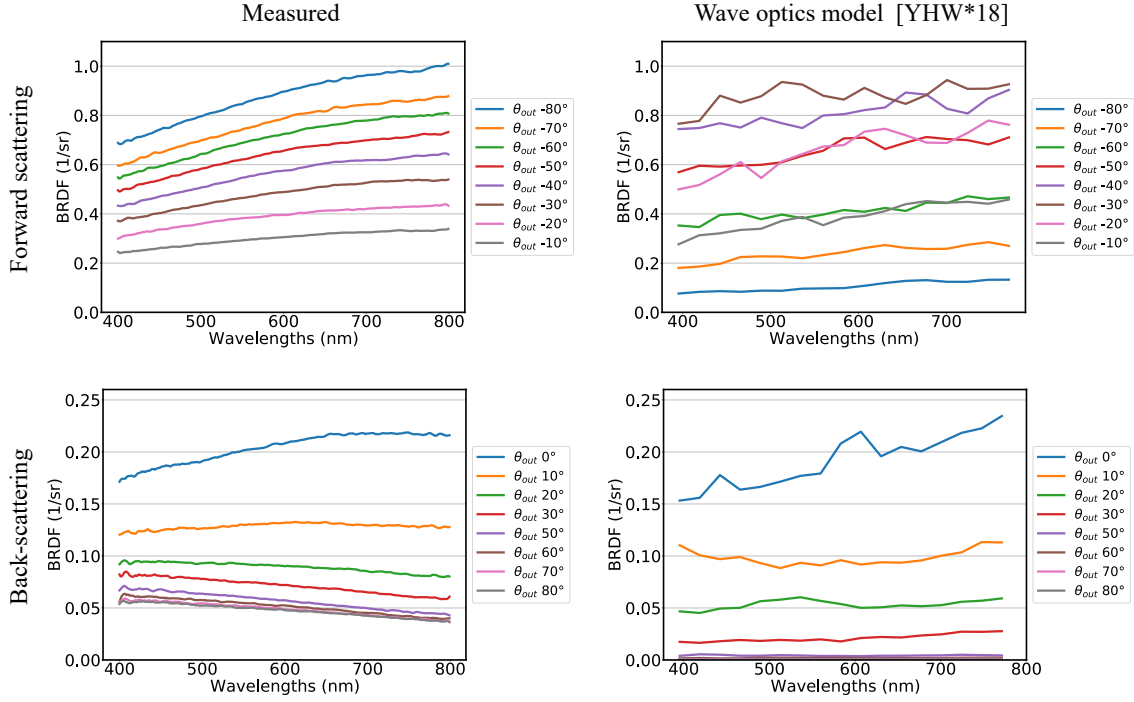


Figure 4.5: On the left side, the measured spectral BRDFs of Sample 1 for forward and back-scattering at $\theta_i = 40^\circ$. On the right, the same for simulated scattering using the wave-optics approach from Yan et al. [3]. The measured BRDF spectra are linearly shifted towards long and short wavelengths in the forward and back-scattering. Even though the increasing trend is noticeable in the simulations, back-scattering is not contemplated.

The spectral data provide further information, as shown on the left side of Figure 4.5. For an incident angle of 40° , we observe close to linear increasing and decreasing spectra for forward and back-scattering, respectively. The spectra within these two extremes are tilted continuously from a positive slope in forward scattering to a negative slope in back-scattering.

When comparing with the Fresnel reflectance of aluminium, we note the general behaviour is present but without the tilt of the spectra. However, it is the only physically plausible wavelength-dependent term in BRDF models based on geometric optics. It is clear that these approaches cannot model the observed colour pattern in the forward scattering and the linear shift tendency.

4.4 Analysis

In the previous section, we introduced three assumptions about the observed wavelength shifts on rough aluminium samples. First, the wavelength shifts originates from diffraction; second, they have a linear behaviour; third, they are shifted towards long and short wavelengths in the forward and back-scattering,

respectively.

These observations are new to the computer graphics and optics community. To the best of our knowledge, only the works by Levesque et al. [48, 49] report these phenomena. Still, none provides a sound investigation or explanation. In the following, we present a thorough analysis of the observed wavelength shifts, providing a better understanding that supports our assumptions and serves as the basis for new and simpler BRDF models.

Diffraction phenomena

To confirm that the observed scattering phenomena are indeed caused by diffraction, we used the BRDF model provided by Yan et al. [3] to simulate the wave-optics scattering of Sample 1. As mentioned in Section 4.1, this model has the advantage of integrating the GHS theory and, consequently, being valid for all angles and rough surfaces. In addition, it works directly on the measured microgeometry.

We first sample the microgeometry in $1\ \mu\text{m}$ steps, and then simulate the in-plane BRDF by evaluating an area of $500 \times 500\ \mu\text{m}$ for 20 different wavelengths.

Figure 4.4 (right) shows the RGB representation of the simulated in-plane BRDF. In the mirror direction, a reddish colour can be observed, similar to the one in the measured data (left). As illustrated on the right side of Figure 4.5, the corresponding simulated spectra have the same increasing trend as the measured data on the left side. The matching reddish colour and the increasing trend indicate that the observed wavelength shift in forward scattering is caused by diffraction.

The simulation, however, does not account for the decreasing trend in back-scattering as observed in the measured data. One reasonable explanation is that multiple reflections are not considered in the simulation. Apart from this, there are other noticeable differences between measured and simulated data.

First, the simulated BRDF has a stronger specular reflection than the measured BRDF, which is illustrated in Figure 4.4 by the narrow forward scattering and the near lack of back-scattering. Again, the ignored multiple reflections are a reasonable explanation.

Second, the simulated BRDF has weaker off-specular peaks than the measured BRDF. Even more, for grazing incident angles, the off-specular peak of the simulated data is shifted towards smaller reflection angles, contradicting the measured data. We noticed similar behaviour with the scatter simulation using the virtual goniometer. Discretisation issues regarding microgeometry are a plausible explanation for this deviation.

Linear wavelength shifts

When dealing with scattering models, a wavelength dependency of $1/\lambda^4$ is often mentioned. A linear wavelength dependency, as observed in Section 4.3, is unusual, which was one of the primary motivations for this research. We further investigated this linear behaviour by approximating the scattering of rough surfaces with the diffraction equation.

The sandblasted aluminium samples have a stochastically created surface with an almost Gaussian height distribution. The superposition of many phase gratings with varying orientations, periods, amplitudes, and phases can approximate this kind of surface [77]. To investigate the observed linear wavelength dependency, we further simplified this approach by approximating the rough surfaces by superpositioning many amplitude gratings with different spacing. For each grating, the relationship between the grating spacing, incident angle, and the angle of the diffracted light is described by the diffraction equation:

$$\theta_o = \arcsin(\sin(\theta_i) - \frac{m\lambda}{d}), \quad (4.1)$$

where the parameters m and d are the diffraction order and grating spacing, respectively. In Figure 4.6 (a), the first order of light diffraction at gratings with varying spacing is plotted. The plot comprises two interesting characteristics. First, when the grating spacing decreases, the light is more strongly diffracted. Second, small wavelengths are diffracted weaker than long wavelengths.

The detector of a gonioreflectometer or the human eye has a limited angular resolution, thus integrating over an angle range. When simulating this integration by binning the grating diffraction with a bin size of 1° , the resulting spectra are shown in Figure 4.6 (b). All spectra have a clear linear dependency, where at diffraction angles of 0° and 1° , the slope is negative and at the remaining angles positive. The negative slope results from the fact that there are practically no long wavelengths in the first two bins.

These results show that the linear dependency does not originate directly from the diffraction on the surface. But instead from the integration of the diffraction over an angle range. Even though our approach is still a rough approximation, it offers a reasonable explanation for the observed linear dependency in the measured data.

Local fit of linear shifts

As stated in Section 4.3, the measured spectra seem to be tilted versions of the Fresnel reflectance spectra. To verify this assumption, we heuristically describe

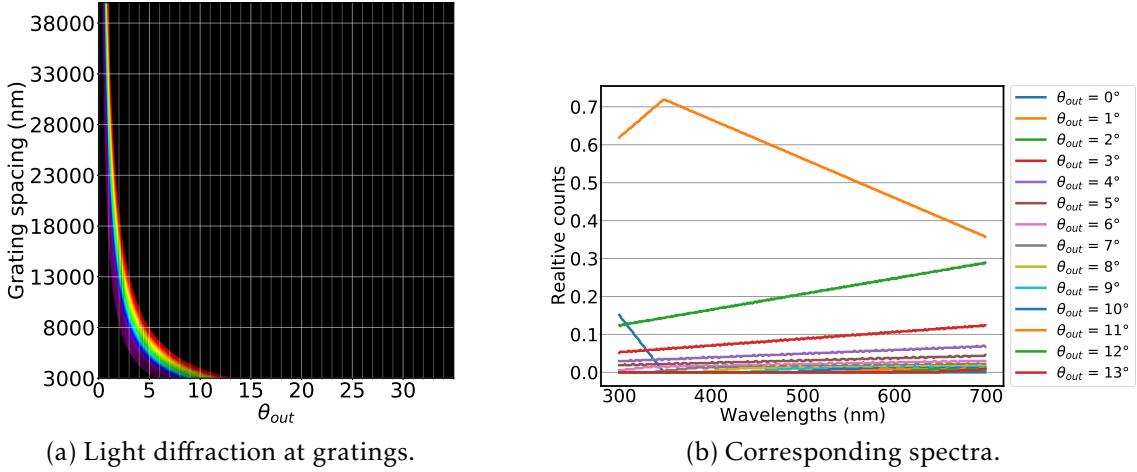


Figure 4.6: First order of light diffraction at gratings with varying spacing (a) and the resulting spectra when binning over the diffraction angles (b). Diffraction is a reasonable explanation for the linear wavelength shifts observed on rough surfaces.

the observed wavelength dependency by the multiplication of the Fresnel term with a linear function:

$$f_{norm}(i, \lambda) = (m\lambda + b) \cdot F^p(i, \lambda), \quad (4.2)$$

where m and b are, respectively, the slope and y-axis intersection of the linear function, and the exponent p takes into account the effect of multiple reflections.

We fit this model against the measured in-plane BRDF of Sample 1. The parameters m , b , and p are determined for each reflectance spectrum separately, and we refer to this approach as local fit. Before conducting the local fit, the BRDF spectra are normalised on an arbitrary wavelength λ_n , in our case 600 nm, to obtain the spectral characteristic and remove any global scaling factor. The Fresnel spectrum is likewise normalised, and b is chosen to rotate the spectrum around λ_n by the linear function. This heuristic model fits the BRDF spectra very well with an r^2 score > 0.95 and > 0.90 for, respectively, 88 % and 93 % of all BRDF spectra.

In Figure 4.7, the fitting results are further investigated from a perceptual perspective. As expected, the Fresnel term has a neutral greyish appearance for almost all incident and reflection angles, while the colour variation comes entirely from the diffraction term. Again, it is evident that the Fresnel term cannot model the colour modulation properly, leading naturally to large ΔE_{00}^* values up to 12. Nevertheless, the Fresnel term still plays an important role, which lowers the ΔE_{00}^* when accounting for it, as in Equation 4.2. In rare cases, barely perceptible differences with ΔE_{00}^* values above the JND of 1 occur.

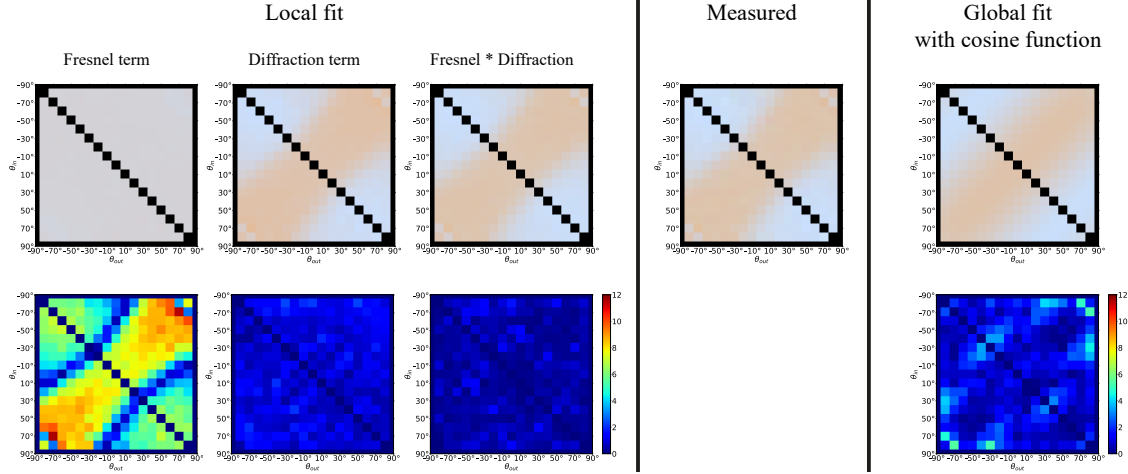


Figure 4.7: Comparison of the local and global fit with the measured data for Sample 1 using the RGB representation. On the left, the local fit is divided into the Fresnel and Diffraction term. The Delta E's 2000 between the fitted and the measured BRDFs are shown on the bottom row. In the case of local fit, the Delta E's remain below the perception threshold. The global fit leads to slightly higher Delta E's but can still adequately reproduce the measured data.

Both the r^2 score and the ΔE_{00}^* values confirm our assumption that the observed diffraction phenomenon has a linear wavelength-dependency, in agreement with the observations made by Levesque et al. [49]. This result is of substantial value for developing a simple BRDF model that considers these phenomena.

Global fit of linear shifts

Albeit the successful fitting by a linear function combined with the Fresnel term, each spectrum was still considered independently. However, a global model is necessary to describe the diffraction phenomena in an analytical BRDF model.

Similarly to the local model, we can describe the global model as follows:

$$f_{norm}(i, o, \lambda) = shift(i, o, \lambda) \cdot F^P(i, \lambda). \quad (4.3)$$

$$shift(i, o, \lambda) = m(i, o) \cdot \lambda + b. \quad (4.4)$$

This model differs from Equation 4.2 only by the dependency of the slope distribution m on the incident and reflection direction. We found that a cosine function parametrised by θ_m describes this distribution well:

$$m_{cosine}(i, o) = h \cdot \cos(w \cdot \theta_m) + t_y, \quad (4.5)$$

where w and h scale the cosine function in the x- and y-axis, and t_y translates the

function along the y -axis. Note that θ_m is the angle between the macro surface normal n and the microsurface normal m .

Figure 4.8 provides an example of a global fit for the slope distributions of Sample 1. For each incident angle, we plot in blue the local fits and in orange the global fit using the cosine function. The figure contains two interesting characteristics. First, the slope values are symmetrical around the mirror reflection, and second, they are periodic. The latter is even more noticeable for smoother samples. We also note that the cosine function fit is adequate, only having issues with the fast transition from a positive to a negative slope. Similar fitting results are achieved for the other samples.

Figure 4.7 shows on the right side the global fit for Sample 1, while Figure 4.9 further shows results for Samples 2-5. Overall, the model provides a good fit with increased precision for rougher surfaces. For smoother surfaces, the colour differences increase, leading to noticeable differences for Sample 5.

The inadequacy of the cosine function to model fast transitions, which are sharper on smoother surfaces, is a reasoned explanation for this trend. Furthermore, for smooth samples, the signal-to-noise ratio is very low for reflection angles away from the specular reflection. This problem is mainly observed in the mirror-like sample, which our model cannot fit properly.

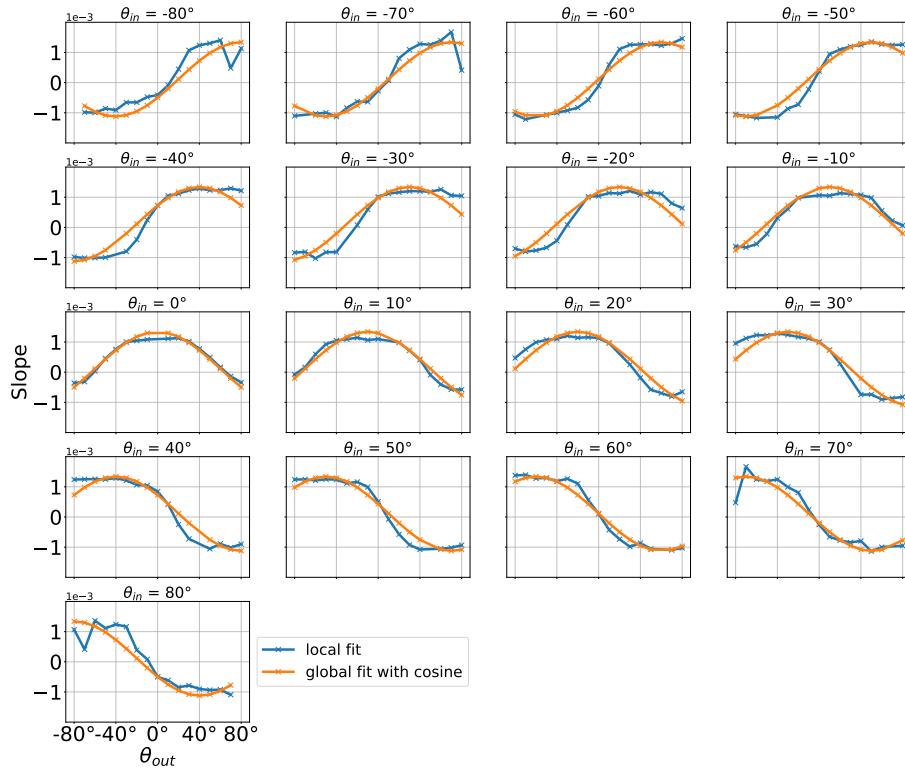


Figure 4.8: Slope distribution (y -axis) of Sample 1 resulting from the local fit (blue line) and the global fit using a cosine function (orange line).

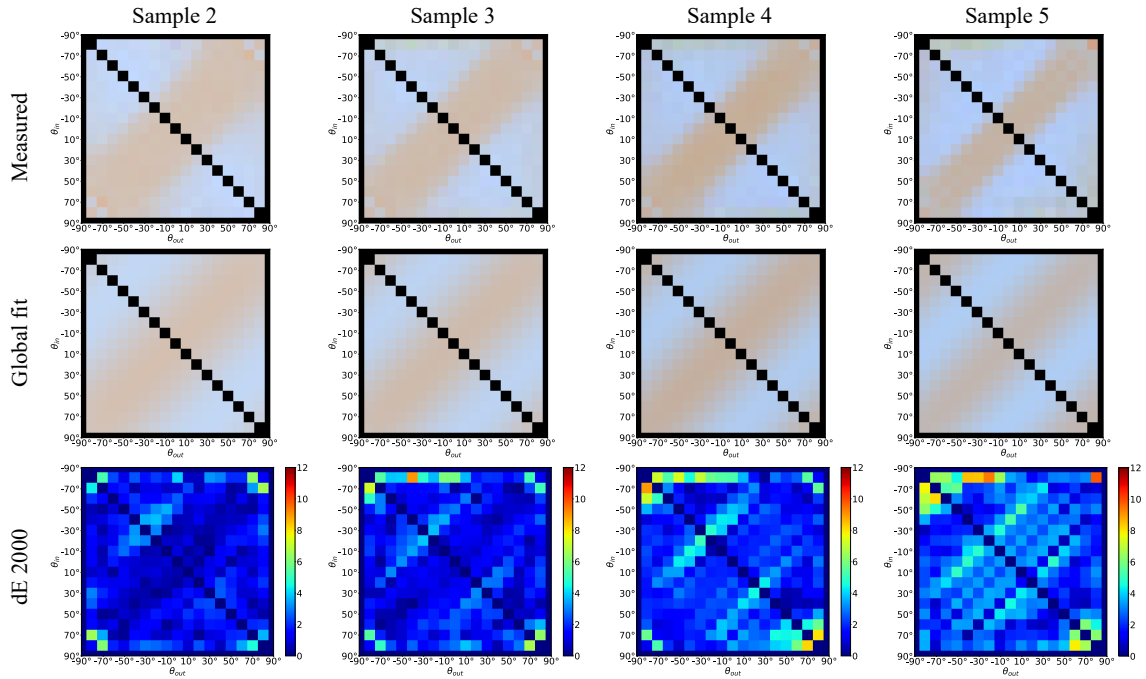


Figure 4.9: Results of our global fitting model for Samples 2-5.

Scatter decomposition

Previously, we focused on analysing light scattering without having information about its composition. It is helpful to decompose the light scattering into the number of reflections to better understand the observed scattering phenomena, mainly the different diffraction phenomena in the forward and back-scattering. This can be achieved using our virtual goniometer (Sec. 2.3) combined with the microgeometries acquired with the confocal microscope. Figure 4.10 shows the decomposition of the light scattering of Sample 1 into the first (blue) and multiple reflections (light blue) for different incident angles. When further decomposing the multiple reflections, we noted that the trend is the same for each number of reflections, so there is no significant difference between two or more reflections in terms of behaviour.

The plot demonstrates that, as expected, first reflections dominate the forward scattering for all incident angles. Away from the forward scattering, the percentage of multiple reflections increases, and, for large incident angles, multiple reflections dominate the back-scattering. There is a correlation between the distribution of the first reflection and the slope distributions in Figure 4.8 confirming our assumption that the wavelength shifts toward long and short wavelengths are related to first reflections in forward scattering and multiple reflections in back-scattering. Furthermore, the scattering simulations explain the plateau observed in the slope distributions, i.e. it is present when first reflections account for nearly 100 % of the captured rays.

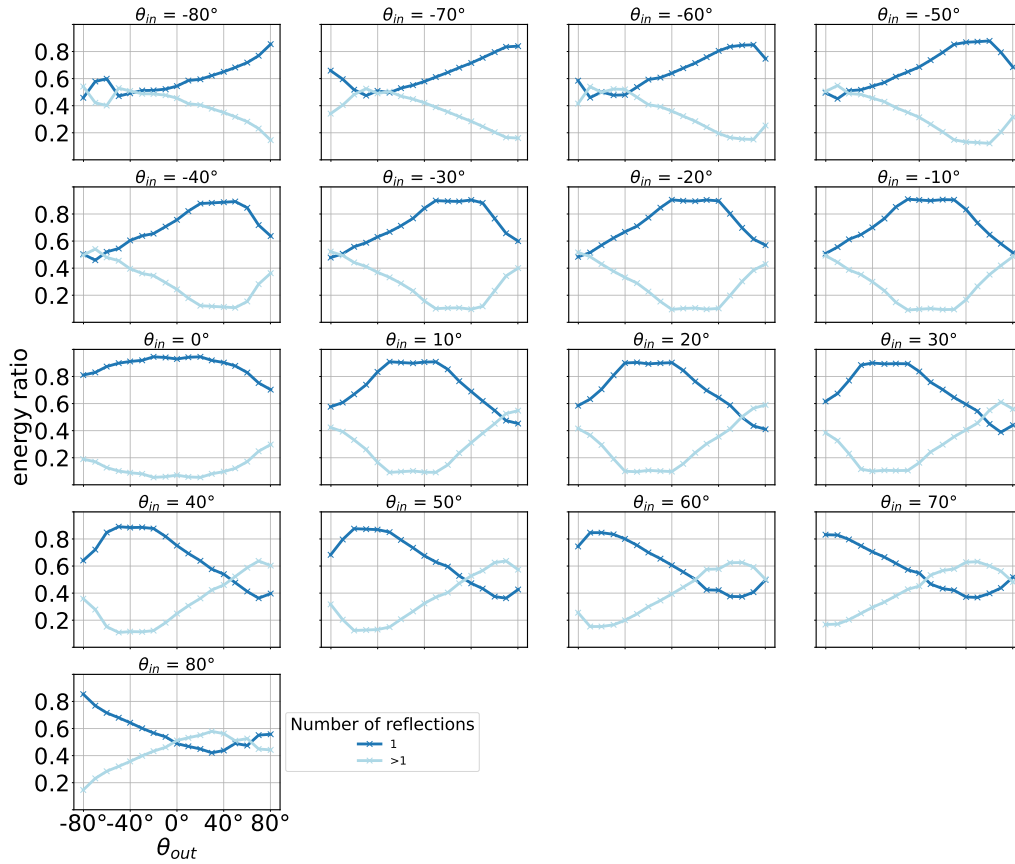



Figure 4.10: Simulation results show in blue the percentage of energy coming from single reflection rays and in light blue multi-reflection rays for wavelength 380 nm and varying incident angles.

The strong influence of multiple reflections in the back-scattering direction also explains why the wave-optics simulation (Sec. 4.4) failed to capture such effects since it does not account for them.


Diffraction on dielectric surfaces

We first observed the wavelength shifts at diffuse, dielectric ColorChecker patches and the grey wall paint (Sec. 3.4). Levesque et al. also reported similar wavelength shifts at Spectralon SRS-99 [49]. The diffraction phenomena are less prominent for these rough dielectrics since internal diffuse scattering dominates the overall reflectance at small angles. Nevertheless, surface scattering increases at grazing angles due to the Fresnel effect, and the shift becomes perceivable.

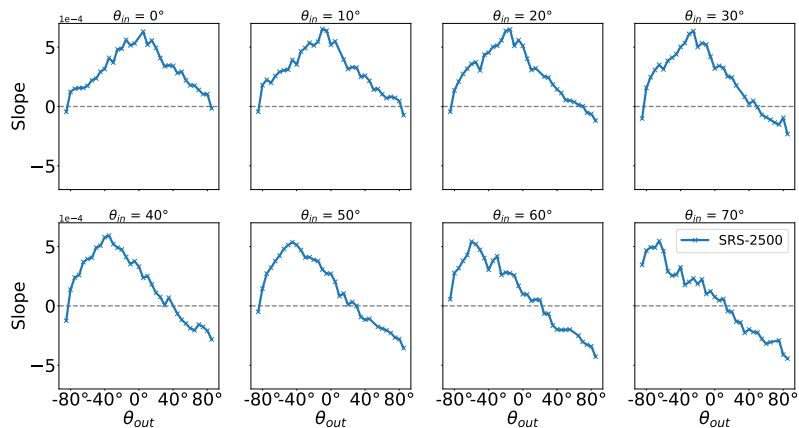
Similar to the aluminium samples, we prepared samples of the Spectralon SRS-02 with different roughness levels. We sanded each sample with sandpaper with varying grain sizes, from 60 to 5000 on the CAMI scale, where lower numbers indicate rougher samples. Then, we measured the in-plane BRDF and retrieved the slope of the spectra with the local fit approach. We assume that SRS-02 reflects light uniformly, so we don't consider the Fresnel effect, which is



(a) SRS-02 60.



(b) SRS-02 1000.



(c) SRS-02 2500.

Figure 4.11: Slope distributions locally fitted with a linear function of Specralon SRS 02 sanded with grain size in the CAMI scale of 60, 1000, 2500. Back-scattering dominates for roughness 60, while forward scattering dominates for roughness 2500. For roughness 1000, forward and back-scattering compensate each other.

equivalent to setting $p = 0$ in Equations 4.2.

Spectralon SRS-02 has a 2% reflectance and scatters light in an almost ideal diffuse manner. Similar to the more well-known Spectralon SRS-99, it has no wavelength preferences. The strong light absorbance eliminates internal diffuse scattering and, thus, allows for the analysis of surface scattering.

When analysing the wavelength shifts of this material, we note some interesting characteristics in line with our previous observations. Figure 4.11 shows the slope distributions for three different roughness (60, 1000, 2500). For roughness 60, back-scattering dominates. At an incident angle 0° , the slopes are distributed as a valley with a maximal negative slope around the incident angle. When the incident angle increases, the valley accompanies the back-scattering direction while the slope for the forward scattering increases. For roughness 2500, we observe a reverse behaviour, which leads us to assume that the forward scattering dominates the back-scattering in this case. Contrary to the previous two roughness, for roughness 1000, the wavelength shifts caused by forward and back-scattering compensate each other, leading to a uniform slope distribution at incident angle 0° .

These results demonstrate that the wavelength shifts occur on conductive and dielectric surfaces, which again confirms the diffraction hypothesis. Furthermore, it confirms our assumption that forward and back-scattering lead to shifts towards long and short wavelengths.

Discussion

From the measured BRDFs, we observed wavelength shifts in the responses, which are not caused by measurement errors, as proven in Chapter 2.2.3. The varied analysis with the acquired multi-modal dataset points towards a diffraction-based effect. We note that the shifts are linear when compensating for the Fresnel effect, an unusual optics behaviour. Our local and global fitting models confirm this linear tendency, and the grating simulations point towards phenomena related to the integration of diffraction effects. We suspected a further dependency on the number of reflections, which the scattering simulations confirmed. Finally, we showed that these are not phenomena particular to conductors by performing the same analysis with a dielectric material.

Based on the presented experimental analysis of light interaction with rough surfaces, we posit the following theory:

Diffraction occurs on rough surfaces, dielectric and conductive, leading to an on average reddish and bluish appearance in forward and back-scattering, respectively. These diffraction phenomena are lin-

early wavelength-dependent; thus, a linear function describes them well. The shift towards long wavelengths originates from light scattered a single time. Conversely, the shift towards short wavelengths originates from multiple reflections. These phenomena dictate the overall visual appearance of rough conductors with a neutral reflectance. On the other hand, at rough dielectrics, they are only prominent at grazing angles.

We showed that the colour shift is, at least for the investigated aluminium samples, adequately fitted by the simple analytic model presented in Equation 4.3. The presented shift function can be used to extend microfacet BRDF models, like the classical Cook Torrance model. As shown in the following section, this is a straightforward and computationally cheap approach to simulate diffraction.

Moreover, the presented model still has limitations as it is only valid for a limited range of surface types. Our model struggles to fit the slope distribution of roughness 1000, where the forward and back-scattering compensate for small incident angles but become more prominent with increasing angles. A more general approach could come from the superposition of two weighted functions separating the diffraction phenomena caused by first and multiple reflections. The weights could be defined by the ratio of first and multiple reflections.

4.5 Application

So far, we analytically described in Equation 4.3 the observed wavelength shifts on rough surfaces caused by diffraction. This section shows how the shift function can easily extend microfacet models based on geometric optics (Eq. 4.4). Further, we demonstrate that the new diffraction model improves the fitting of measured BRDFs of conductors and dielectrics in all cases. Finally, we compare renderings computed with the Cook-Torrance and our diffraction model.

BRDF model

As stated in Chapter 2.1.4, the Cook-Torrance model is a widespread BRDF model based on microfacet theory that models the reflection of dielectrics and conductors. In the following, we will use two slightly modified versions of the original Cook-Torrance model, one for conductors and one for dielectrics. The reflection of conductors is described by the microfacet model as follows:

$$f_r(\lambda, \omega_o, \omega_i) = \frac{D(\omega_h)F(\omega_i, \omega_h)G(\omega_o, \omega_i, \omega_h)}{4 \cos(\theta_o) \cos(\theta_i)}. \quad (4.6)$$

In contrast to the original Cook-Torrance model, the diffuse and specular albedo are removed because conductors don't refract light; hence they don't re-emit light, and the specular albedo isn't physically plausible.

For dielectrics, the model for conductors is extended by the diffuse albedo to account for internal scattering:

$$f_{dielectric}(\lambda, \omega_o, \omega_i) = \frac{k_{d,\lambda}}{\pi} + f_r. \quad (4.7)$$

In the following, we will refer to both models as modified Cook-Torrance model.

The Cook-Torrance model is based on geometric optics; hence wave-optics phenomena, such as the observed wavelength shifts, are not yet considered. This leads, particularly for rough conductors, to clear perceivable deviations from reality. Therefore, we extended the modified Cook-Torrance models with the observed wavelength shifts by multiplying f_r with the shift function as described in Equation 4.4:

$$f_{r,diffraction}(i, o, n, \lambda) = shift \cdot f_r. \quad (4.8)$$

Our model assumes that the overall scattering is well described by the Cook-Torrance model, consequently, by geometric optics and that the shift function approximates the observed diffraction phenomena.

We observed that for both the conductive and dielectric versions of our BRDF model, the exponent p and the offset parameter t_y of the global fit approach (Eq. 4.3) have a marginal influence on the fitting of the measured BRDF. Hence, we did not include them in the BRDF model.

Fitting

A common way to validate BRDF models is to fit them against measured BRDF data. To demonstrate that our diffraction model outperforms the modified Cook-Torrance model, we fit both against the measured in-plane BRDFs of the aluminium samples (conductors) and the 24 ColorChecker patches (dielectrics) introduced in the previous chapter. The GGX distribution function and Smith's shadowing-masking function, described in Chapter 2.1.4, are used for all models. The implementation of the Fresnel equation and the fitting process is different for conductors and dielectrics, as explained in the following.

Aluminum samples: To fit the in-plane BRDFs of the aluminium samples, we use the original Fresnel equations as defined in Chapter 2.1.3 with the refractive index of aluminium given in the literature [78]. Consequently, for the modified Cook-Torrance model, only the roughness parameter α and for our model addi-

tionally, the parameters w and h of the shift function have to be determined.

The fitting of the parameters is performed in python with the nonlinear least squares minimisation function of the library *lmfit* [82]. The Levenberg-Marquardt fitting method and the cost function by Löw et al. [36] are used.


The fitting result of Sample 1 is illustrated in Figure 4.12a. In the top row, the fitted in-plane BRDF with the modified Cook-Torrance model (left) and our model (right), as well as the measured in-plane BRDF (centre), are shown. The bottom row shows the ΔE_{00}^* between the respective fitted and measured BRDF. The figure demonstrates that the scatter distribution of the modified Cook-Torrance, and our model are the same, but the colour appearance is different. Our model fits much better the yellowish and bluish appearance of the measured BRDF in the forward and back-scattering. At all measurement points, our model outperforms the modified Cook-Torrance model as confirmed by the ΔE_{00}^* plots and the, on average, lower ΔE_{00}^* by 2 points. Similar fitting results can be observed in all sandblasted aluminium samples as shown in Table 4.2. The fitting improvements diminish with decreasing roughness, which can be explained by the diminishing influence of the wavelength shifts on smooth surfaces.

Although our model improves the fitting of the measured BRDFs, there are still differences. Our model struggles to fit the scattering distribution while fitting the colour appearance well. We assume that the micro surface of the aluminium samples, which consists of randomly distributed cavities, is not well represented by microfacet models, which approximate the micro surface by tiny flat mirrors. It is well known that not only the distribution of the microfacet normals but also their arrangement have a significant influence on the scatter distribution.


For Sample 6 (mirror sample), we observed another fitting issue. The measured BRDF values at grazing incident angles were much smaller than the simulated ones. A possible explanation is that our goniometer fails to correctly measure the BRDF of smooth surfaces under grazing angles, but more investigation is necessary. In any case, for this sample, measurements under grazing incident angles are weighted less in the fitting process.

ColorChecker patches: The ColorChecker patches' fitting differs from the aluminium samples' fitting due to the unknown refractive index and the diffuse albedo. Both have to be determined in addition to the fitting parameters of conductors.

We use the convenient approximation of the Fresnel equation for dielectrics introduced by Cook and Torrance [29] as reviewed in Chapter 2.1.4. Assuming that the refractive index of dielectrics is wavelength-independent further simplifies the Fresnel equation.



(a) Aluminum Sample 1.



(b) ColorChecker Patch 4.

Figure 4.12: Comparison of the fitting of a conductor (a) and a dielectric (b) with the Cook-Torrance and our model.

Like the fitting process described in Chapter 3.3.2, the process is again divided into a nonlinear and linear optimisation process. The model parameters

Al Sample No.	1	2	3	4	5	6
Cook-Torrance GGX	6.94	7.11	7.74	9.23	6.97	17.48
Our	5.02	5.83	6.32	7.46	6.02	17.05

ColorChecker No.	1	4	10	14	22	24
Cook-Torrance GGX	1.69	1.82	1.47	1.23	1.66	2.37
Our	0.95	0.94	1.09	0.81	1.27	0.92

Table 4.2: Averaged dE 2000 between the measured in-plane BRDFs of the aluminium samples (top table) and ColorChecker patches (bottom table) and their respective fittings with the Cook-Torrance and our model. For all samples, our model outperforms the Cook-Torrance model.

α , η , w and h are determined in the nonlinear optimisation, where w and h are only determined for our diffraction model. For the linear optimisation, linear regression is employed to determine the wavelength-dependent diffuse albedo $k_{d,\lambda}$.

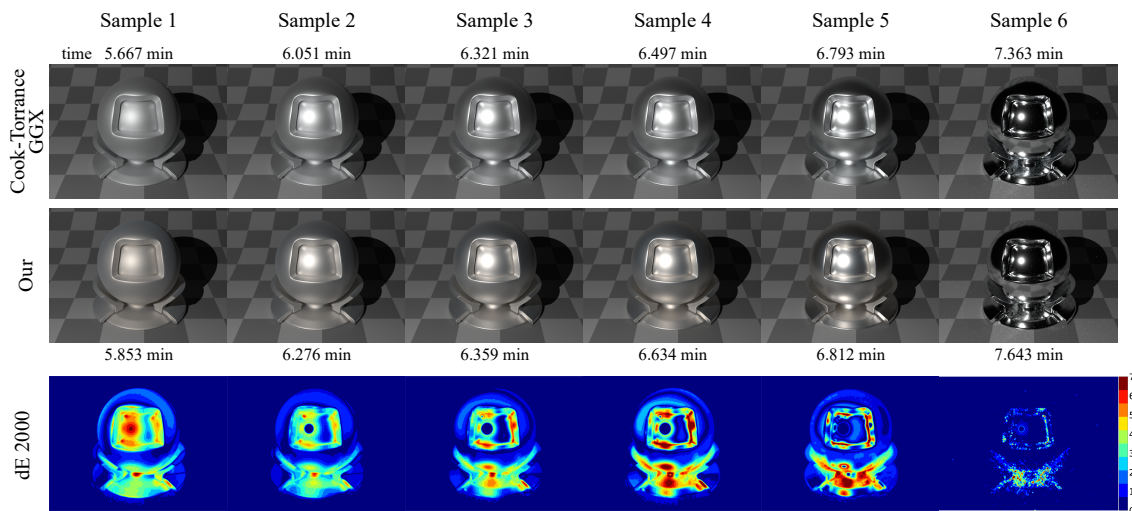
The fitting result of the ColorChecker Patch 4 in Figure 4.12b demonstrates that our model outperforms the modified Cook-Torrance model also when fitting dielectrics. Our model improves the fitting of diffuse reflection and of the specular reflection, particularly visible at grazing angles. The modified Cook-Torrance model tries to compensate for the lack of wavelength shifts with the diffuse albedo resulting in a reddish-tinted diffuse albedo. The ΔE_{00}^* plots confirm this observation; the ΔE_{00}^* is lower at grazing and small angles, where the diffuse reflection dominates.

Similar fitting results are obtained for all ColorChecker patches. The results for six patches are shown in Table 4.2.


Renderings

The fitting results demonstrate that our model improves the modelling of both conductors and dielectrics. Nevertheless, it provides limited information on the visual differences of objects rendered with the measured, the modified Cook-Torrance, or our model. Although we cannot generate reference renderings with the measured in-plane BRDFs, as is usually done, we generated renderings with the modified Cook-Torrance and our model of the aluminium samples and ColorChecker patches. Comparing both renderings demonstrates the influence of the wavelength shifts on the object’s appearance.

We used Mitsuba 2 [4] and extended the Mitsuba plugins *roughconductor* and *roughdielectric* by the shift function (Eq. 4.8). In the case of dielectrics, we used a modified version of the *matpreview* scene, where the environment map is replaced by a spotlight and another directional light source.



(a) Aluminum samples.



(b) ColorChecker patches.

Figure 4.13: Renderings computed with Mitsuba 2 [4] and the fitted BRDFs of aluminium samples (a) and ColorChecker patches (b). Clear differences between the Cook-Torrance and our model can be observed for both conductors and dielectrics, as confirmed by large Delta E 2000 values. Note the different scales of the colour bars.

In Figure 4.13, the renderings of six aluminium samples (a) and six ColorChecker patches (b), the corresponding render times, and the ΔE_{00}^* images are presented. The modified Cook-Torrance model creates a homogeneous, greyish appearance for the aluminium samples. On the contrary, with our model, the object is tinted yellowish and, towards the edges, bluish. This effect can be well observed on rough surfaces, which is confirmed by the higher ΔE_{00}^* values. In our opinion, our model leads to a more natural, realistic look, while the object appears more synthetic with the Cook-Torrance model.

For the ColorChecker patches, the differences between the renderings are more subtle, leading to smaller ΔE_{00}^* than for the aluminium samples. The differences are difficult to perceive in the side-by-side comparison. However, the ΔE_{00}^* are nearly for all patches above the JND of 1 and for some, particularly dark patches, clearly above the JND. These differences probably further increase for dielectrics with lower refractive indices due to the stronger surface reflection.

The render time of the modified Cook-Torrance and our model is practically the same. Only for conductors an increase of up to 3 % is observed. The simplicity of the model, accompanied by the low computational costs, are huge benefits compared to current diffraction models and, consequently, enable its use in real-time applications.

4.6 Conclusion

Our experimental investigation of light interaction with rough surfaces not only gives a better understanding of scattering phenomena but also allows us to describe them in a simple way. We combined different measurement modalities to thoroughly analyse some underlying phenomena that significantly influence the material's appearance.

Based on the analysis of the experimental results, we posited the theory that single reflections, predominately in the forward scattering, leading to a linear wavelength shift towards long wavelengths, thus to a reddish appearance. Contrary, multiple scattering, predominately in back-scattering, leads to a linear shift towards short wavelengths, hence a bluish appearance. We provide strong evidence that both phenomena are caused by diffraction. We observed that these linear wavelength shifts significantly impact the appearance of rough conductors. It also influences, even if less prominently, the appearance of dielectrics. Our conclusions are in agreement with all different measured data modalities. Furthermore, simulations using wave and geometric optics back our claims.

These are novel insights for the computer graphics community, and breaking down the scattering phenomena allows for considering diffraction effects in

a simplified manner. This is well demonstrated by the introduced diffraction model, where the Cook-Torrance model is extended by the shift function to account for diffraction. The simple approach outperforms the Cook-Torrance model for all tested dielectrics and conductors at nearly no extra computational costs. The provided renderings further demonstrate the high impact of the wavelength shifts on the visual appearance.

5. Colour in virtual reality

As we motivated in the introduction, the image display is, besides the light simulation, another bottleneck of VP. To further improve the fidelity of VP and thus increase the acceptance in the industry, we address in this chapter the colour-consistent image display on VR HMD.

Usually, rendered images are displayed on a monitor. Current monitors provide consistent colour reproduction with a large colour gamut, high resolution and high dynamic range. However, conventional monitors are rarely used for VP due to their low immersion. Virtual prototypes require complex visualization systems, providing stereo reproduction with head tracking, a display with high resolution and large field of view (FOV), and interaction possibilities.

Current systems addressing these properties are often based on virtual reality (VR) or augmented reality (AR). VR describes a computer-generated virtual world where users can move freely and interact with virtual objects in real-time. The primary goal of VR is that the user dives into the virtual world and loses touch with reality. The virtual world has not necessarily been similar to the real world.

AR describes the enhancement of the real world by computer-generated content, usually three-dimensional objects, integrated into the real environment. Therefore, AR devices have to scan the real world geometrically. Visualization systems based on AR usually provide poor colour rendering due to the physical superposition of the virtual content with the real environment. A consistent colour rendering compensating the background is complicated and reduces the colourspace considerably. Consequently, virtual prototypes used for design decisions often use visualization systems based on VR.

In the past, Powerwalls and CAVEs were predominately used as VR systems for VP. The Powerwall is a large display with a high resolution, often consisting of several displays or projectors arranged in a matrix. The CAVE follows a similar approach; however, the user is surrounded by three up to six projection walls arranged in a cube instead of a single wall. Although both systems provide stereo reproduction and head tracking, the immersion of a CAVE is much better due to the considerably larger FOV. The major disadvantages of these VR systems are

the high price and the required physical space. Furthermore, in both systems, immersion is often disturbed by the mixture of virtual and real environments.

As stated in Chapter 1, virtual reality head mounted displays (VR HMDs) have become more attractive to the industry with the advent of recent consumer devices, such as the Oculus Rift and the HTC Vive. They provide a similar or even higher immersion as CAVEs [83] or Powerwalls while being significantly cheaper and requiring less space. In the industry, consumer VR HMDs are already in use for VP, although there has been no work dealing with the display characterization and calibration of VR HMDs from a colour perspective. In this chapter, we demonstrate that colour calibration is crucial for the fidelity of virtual prototypes. We present a simple display model for VR HMDs and verify the simulated and real application of the display model using two frameworks. The display model and the proposed frameworks are implemented with the HTC Vive Pro and the Pimax 5k+.

5.1 Related work

Only a few works deal with the colour calibration of VR systems, and none consider the colour calibration of VR HMDs. Some works describe the characterization and calibration of a stereoscopic projection system based on the Infitec colour separation technology [84–86]. The Infitec technology requires two projectors with slightly different primary valences, often achieved by using interference filters in front of the projectors. The user wears special glasses with the corresponding filters, which separate the full-colour image for the left and right eye [87].

Kresse et al. [84] built two multi-projector displays using the Infitec technology; the digital CAVE and the HEyeWall (a high-quality and high-resolution stereo display). To provide colour consistency, they determined the common gamut of all projectors and the gamma curves for each colour channel. Adjusting the input-RGB values accordingly improves the colour consistency but still leads to clearly discernible differences between different projectors. Furthermore, using a common gamut decreases the gamut and contrast compared to the native properties of the projectors.

Gadia et al. [85] present a virtual reality theatre which uses a stereo multi-projector display similar to the HEyeWall. They characterize the VR theater’s projectors spectrally and colourimetrically. They conclude that the filters reduce the maximum luminance from 44cd/m^2 to 15cd/m^2 and that the common gamut negatively affects the red and blue channels.

Gerhardt et al. [86] extend the approach of Kresse et al. [84] by adding a

colour difference threshold. The common gamut is maximized in an optimization process where the colour differences are within the defined threshold. They recommend a threshold of $\Delta E_{a^*b^*} = 20$ for a white uniform patch. This threshold is clearly above the JND of 2.7.

Projector-based VR systems are costly and require a lot of space. It is difficult to calibrate them, and even calibrated systems are insufficient for colour-critical applications such as VP. Recently, Infitec published new filters with more spectral bands. Referring to their website [88], these filters lead to more accurate colour reproduction. Unfortunately, no evaluation results are available yet.

5.2 Definition and verification of the display model

There are two classical ways to calibrate displays: with a look-up table containing many display measurements or with a matrix profile. The latter is typically preferred due to its simplicity. The matrix approach requires a display model that describes how a digital input signal is transformed into a specific rendered colour on display. In the following, we present a display model for the calibration of VR HMDs and two frameworks for verifying how well displays respect the assumptions made in the display model.

The display model

We use the display model shown in Figure 5.1 to describe the colour reproduction of VR HMDs. Our model assumes that each pixel consists of three independent light sources with the constant emission spectra $r(\lambda)$, $g(\lambda)$ and $b(\lambda)$. Their intensities are controlled by the linear gain factors RGB_{Lin} within the range $[0, 1]$. The gamma curves describe the nonlinear relationship between the RGB_{Lin} and the input signals RGB . The three additive primary colours are overlaid by an in-

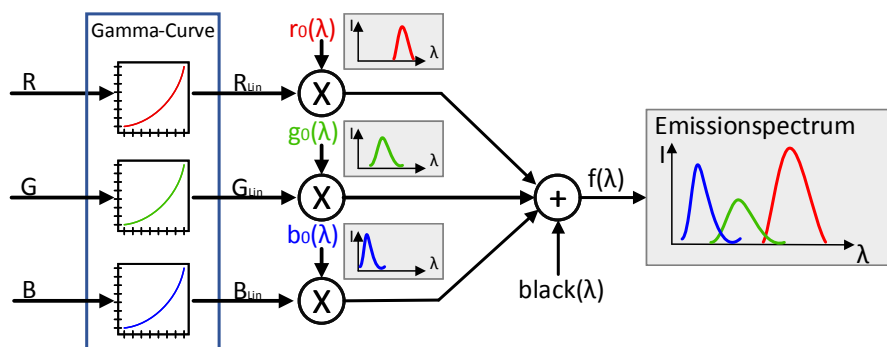


Figure 5.1: The Display model.

dependent black component comprising the residual display light and ambient light. For HMDs, the reflection component is negligible due to the encapsulation of the displays from ambient light.

According to the display model, the colour stimulus $f(\lambda)$ is linearly composed of the intensities of the modulated three primary spectra and the black spectrum, as shown in Equation 5.1:

$$f(\lambda) = black(\lambda) + R_{Lin} \cdot r_0(\lambda) + G_{Lin} \cdot g_0(\lambda) + B_{Lin} \cdot b_0(\lambda) \quad (5.1)$$

Based on this equation, a relationship between the RGB_{Lin} signal and the resulting display tristimulus values XYZ_{Disp} is derived in the following.

The XYZ_{Disp} values are computed by multiplying and integrating the colour stimulus $f(\lambda)$ with the CIE colour matching functions (Sec. 2.1.2):

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix}_{Disp} = \int_{\lambda} \begin{pmatrix} \bar{x}(\lambda) \\ \bar{y}(\lambda) \\ \bar{z}(\lambda) \end{pmatrix} \cdot f(\lambda) d\lambda \quad (5.2)$$

Substituting $f(\lambda)$ from Equation 5.1 into Equation 5.2 leads to Equation 5.3, which can then be resolved in terms of the RGB_{Lin} signals as detailed in Equation 5.4.

$$\begin{aligned} \int_{\lambda} \begin{pmatrix} \bar{x}(\lambda) \\ \bar{y}(\lambda) \\ \bar{z}(\lambda) \end{pmatrix} \cdot black(\lambda) d\lambda + \begin{pmatrix} \bar{x}(\lambda) \\ \bar{y}(\lambda) \\ \bar{z}(\lambda) \end{pmatrix} \cdot R_{Lin} \cdot r_0(\lambda) d\lambda + \begin{pmatrix} \bar{x}(\lambda) \\ \bar{y}(\lambda) \\ \bar{z}(\lambda) \end{pmatrix} \cdot G_{Lin} \cdot g_0(\lambda) d\lambda + \begin{pmatrix} \bar{x}(\lambda) \\ \bar{y}(\lambda) \\ \bar{z}(\lambda) \end{pmatrix} \cdot B_{Lin} \cdot b_0(\lambda) d\lambda \\ = \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}_{black} + R_{Lin} \cdot \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}_{Red} + G_{Lin} \cdot \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}_{Green} + B_{Lin} \cdot \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}_{Blue} \end{aligned} \quad (5.3)$$

$$\begin{pmatrix} R_{Lin} \\ G_{Lin} \\ B_{Lin} \end{pmatrix} = \begin{pmatrix} X_r & X_g & X_b \\ Y_r & Y_g & Y_b \\ Z_r & Z_g & Z_b \end{pmatrix}^{-1} \cdot \left(\begin{pmatrix} X \\ Y \\ Z \end{pmatrix}_{Disp} - \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}_{black} \right) \quad (5.4)$$

Finally, we compute the nonlinear RGB control signals by applying the inverse gamma curves. In conclusion, to fully describe our display model, we need to acquire the primary valences, the gamma curves, and the black signal.

Framework 1: Verification of the display model

In the display model, we made several assumptions about the display colour reproduction, which are not always respected by displays. Particularly, VR-HMDs' displays are not well studied in the literature, which is why it is essential to verify

how well they respect the model assumptions. Therefore we provide two verification frameworks, where the first one is depicted in Figure 5.2. The framework is divided into two paths: the measurement and the simulation path. In the measurement path, RGB test colours are displayed, and the resultant emission spectra are measured with an EyeOne Pro 2 (Sec. 2.3). The measured spectra are first converted into tristimulus values and then into the $L^*a^*b^*$ -space.

In the simulation path, the display model simulates the emission spectra resulting from the RGB test colours, which are again converted into the $L^*a^*b^*$ -space. The display model is verified by conducting a physical and colourimetric comparison. In the physical comparison, the measured and simulated spectra are compared, while the ΔE_{00}^* is computed in the colourimetric comparison.

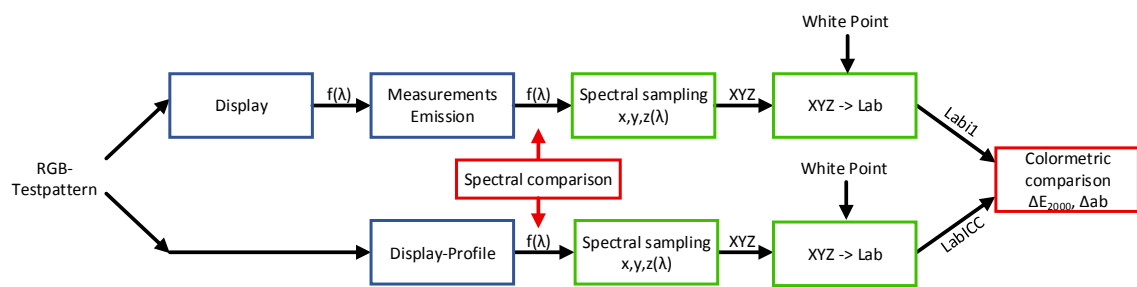


Figure 5.2: Verification of the display model.

Framework 2: Verification of the display model application

The second framework illustrated in Figure 5.3 verifies the application of the display model by calculating the ΔE_{00}^* between measured colours and their reproduction on display. Therefore, the reflectance spectra of a ColorChecker are measured, multiplied with an illumination spectrum, and converted into the absolute tristimulus values XYZ_{Ref} . The XYZ_{Ref} values are normalized with the display black and white point as described in [89]. Applying the display model

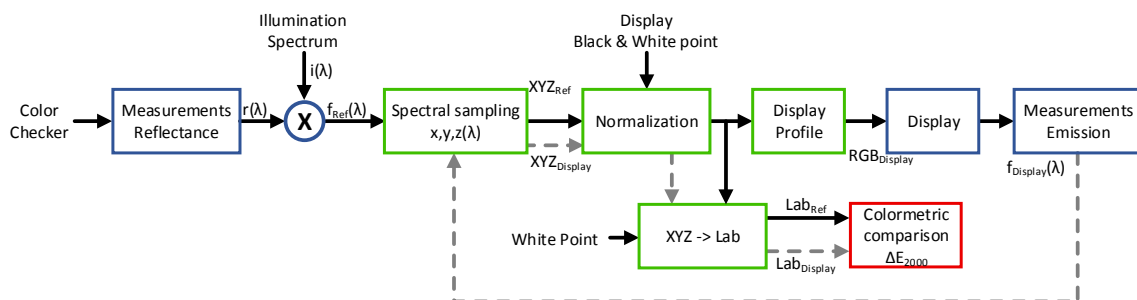


Figure 5.3: Verification of the display model application.

backward on the normalized XYZ_{Ref} values leads to the display control values RGB_{Disp} , which are displayed, and the resulting emission spectra $f_{Display}(\lambda)$ captured with an EyeOne Pro 2. These emission spectra are converted to the tristimulus values XYZ_{Disp} and normalized in the same way as XYZ_{Ref} . Both the normalized XYZ_{Ref} and $XYZ_{Display}$ values are converted with a given white point to $L^* a^* b^*$ -space, and the ΔE_{00}^* is calculated. The resulting ΔE_{00}^* define how well the display model works in the application.

5.3 Calibration of HTC Vive Pro and Pimax 5k+

To verify how well consumer VR HMDs respect the assumptions made in the display model, we implement the described display model and verification frameworks with the consumer VR HMDs: HTC Vive Pro and Pimax 5k+. The HTC Vive Pro is a widespread VR HMD that has already been used successfully in the industry and science. The Pimax 5k+ has an extraordinary display resolution, a weakness of most VR HMDs.

Implementation of the display model

Usually, VR HMDs have two separate displays for the right and left eye. Each display has to be calibrated individually since the display characteristic varies from panel to panel. Even if only one display is used for both eyes, a separate calibration is recommended to compensate for inhomogeneities.

We implement the display model by acquiring the gamma curves, primary valences, and black- and white levels of the left and right display of the Vive Pro and Pimax. For this purpose, each display is controlled with 10 RGB values per colour channel, giving the nonlinear gamma curves and primary valences, as well as with the maximum and minimum RGB signal providing the black- and white level. In all cases, the emission spectra are measured with an EyeOne Pro 2.

The display model defines the device colourspace and, consequently, provides information about the display quality from a colour perspective. A comparison of the display properties of the Vive Pro and Pimax is shown in Figure 5.4, where the diagrams in (a)-(c) characterize the display colourspaces and in (d) verify the display model.

Figure 5.4(a) demonstrates that the Vive Pro has a much lower black level than the Pimax, which can be traced back to the different display technologies. The Vive's AMOLED display completely turns off each pixel independently, whereas the Pimax's LCD panel can not fully block the backlight. This results in a lower and consequently better black level of 0.04 cd/m^2 for the Vive Pro compared to

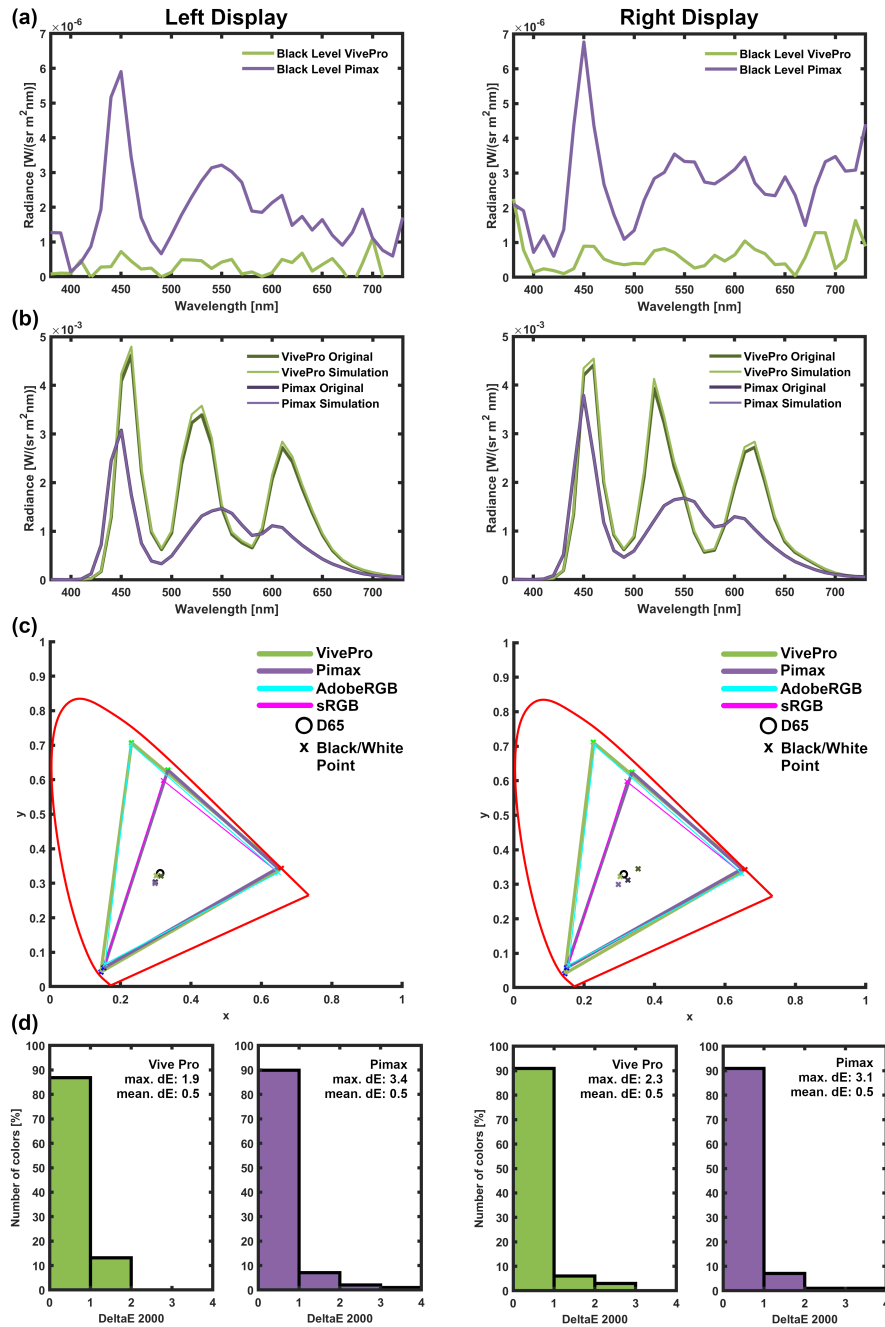


Figure 5.4: Display characterization of the HTC Vive Pro and Pimax 5k+. The left and right columns show the results for the left and right displays. (a) Black level. (b) Original white spectrum and the simulation with the display model. (c) Chromaticity diagram including the sRGB, AdobeRGB, and display colourspaces. (d) DeltaE histogram (see Figure 5.2).

0.21 cd/m^2 for the Pimax.

In Figure 5.4(b), the measured white spectrum (original) and its simulation with the display model are shown. The Vive Pro has a considerable brighter white spectrum than the Pimax, which is confirmed by the maximum luminance of

130 cd/m^2 compared to 80 cd/m^2 . However, the Pimax shows a better match between the original and simulated white spectrum, which indicates that the Vive Pro does not entirely meet the assumptions made in the display model.

The chromaticity diagrams in Figure 5.4 (c) depict the primary valences of the display, sRGB, and AdobeRGB colourspaces. The primary valences of the Vive Pro and Pimax are in great accordance with the AdobeRGB and sRGB colourspace, respectively.

We conclude that the Vive Pro has a considerably larger colourspace than the Pimax. The Vive Pro almost covers the whole AdobeRGB colourspace. Only the maximum luminance of 130 cd/m^2 is below the 160 cd/m^2 of the AdobeRGB colourspace. The Pimax covers only the clearly smaller sRGB colourspace.

Verification of the display model

The Vive Pro clearly outperforms the Pimax in terms of display capacities. However, the display capacities do not provide information about the consistent colour reproduction. To evaluate how well both HMDs respect the model assumptions, we implement the verification framework in Figure 5.2. We used equally sampled RGB values with a step size of 85 digits as the input RGB-test colours. The histograms in Figure 5.4(d) show the resulting ΔE_{00}^* between the measured and simulated test colours. Both the Vive Pro and the Pimax have an average ΔE_{00}^* of 0.5, which is below the JND. Approximately 10% of the test colours lead to ΔE_{00}^* larger than 1, but only a few of them are larger than 2. Hence, most of the colour differences are not discernible when compared side-by-side. The verification of the display model shows that both VR HMDs are in great accordance with the assumptions made in the display model. Thus, the presented display model adequately describes the displays of these two devices.

Table 5.1 summarizes the results of the display characterization and calibration of the Vive Pro and Pimax. From the table follows that the Vive Pro is more suitable for VP than the Pimax due to its larger colourspace, lower black level and lower maximum ΔE_{00}^* .

Verification of the display model application

We finally investigate the display model application by implementing the second framework in Figure 5.3. Following, we only present the results of the left display of the Vive Pro since the left and right displays have almost the same characteristics, and the Vive Pro is more suitable for VP than the Pimax.

As input for the framework, we use the reflectance spectra of the X-rite ColorChecker Digital SG with 96 patches. We virtually illuminate the patches with

	HTC Vive Pro	Pimax 5k+
Display technology	AMOLED	LCD
FOV	110°	200°
Resolution	1440 x 1600	2560 x 1440
Colourspace	~ AdobeRGB	~ sRGB
Max. Luminance	130 cd/m^2	80 cd/m^2
Min. Luminance	0.04 cd/m^2	0.21 cd/m^2
Contrast ratio	3250:1	377:1
Max. dE	2.3	3.4
Mean dE	0.5	0.5

Table 5.1: Results of our display characterization

the display’s whitepoint to guarantee that all ColorChecker spectra are within the display’s colourspace. Usually, VR HMDs are not calibrated, and the HMDs are controlled directly by the sRGB or AdobeRGB values. To demonstrate the differences between a workflow with and without colour management, we conduct the framework with the sRGB, AdobeRGB and measured display profiles.

The results of the three approaches are shown in Figure 5.5, where the first row shows the colour comparison between the reference colours (outer square) and the displayed colours (inner square), and the second row the corresponding ΔE_{00}^* histograms. The figure demonstrates impressively that the standard workflow with the sRGB profile leads to clearly perceptible colour differences between the reference and displayed colour for all patches. The AdobeRGB colourspace considerably improves colour reproduction due to its similarity to the Vive Pro colourspace. However, the colour reproduction is with an average ΔE_{00}^* of 5.7,

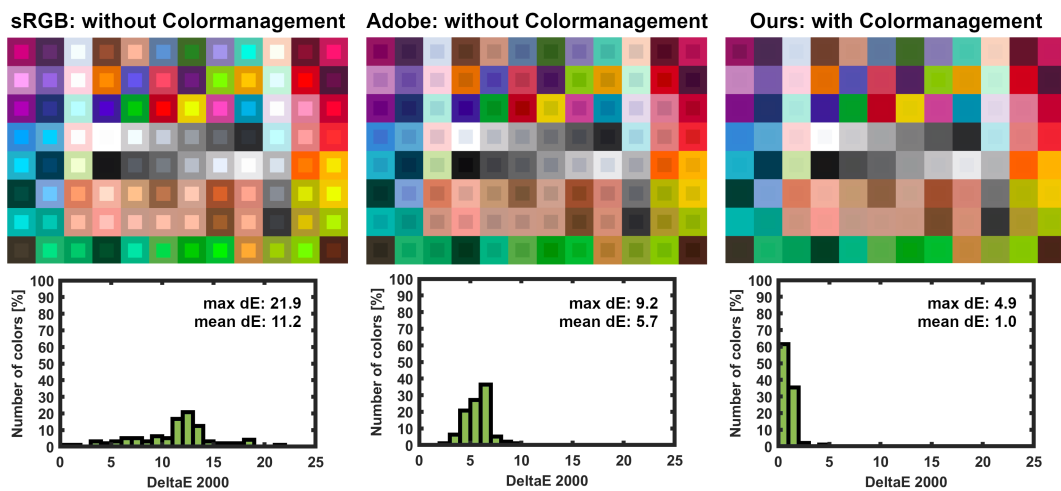


Figure 5.5: Comparison of measured colours and its reproduction with the sRGB, the AdobeRGB, and the measured display profile on the Vive Pro.

still insufficient for colour-critical application. Contrary, our approach leads to a convincing colour reproduction with an average ΔE_{00}^* of 1. The small colour differences again confirm that our display model works and that the Vive Pro, combined with our colour management, can be used for colour-critical applications, such as VP.

5.4 Conclusion

We presented a simple display model, which can be used to characterize the displays of current consumer VR HMDs. We concluded that both the HTC Vive Pro and Pimax 5k+ match the assumptions made in the display model well. However, we recommend for colour-critical applications the Vive Pro due to its considerably larger colourspace. The Vive Pro almost covers the AdobeRGB colourspace, while the Pimax only covers the sRGB colourspace. We demonstrate that with our calibration approach, measured colours can be, on average, reproduced with a ΔE_{00}^* of 1, which is below the JND. The commonly used sRGB and AdobeRGB profiles without colour management lead to average ΔE_{00}^* of 11.2 and 5.7, which is insufficient for colour-critical applications. Another advantage of making our display model suitable for VP is that it can be easily incorporated into a real-time rendering system with only marginal computational overhead.

Although our display model considerably improves the colour reproduction of VR HMDs, there is still space for improvement. Our display model does not consider the Fresnel lenses, thus omitting lens aberrations. The display model has to be extended by the lens's characterization to provide a more complete model. Maxwell et al. [90] already proposed some research in this direction. Furthermore, our display model assumes that the brightness and colour reproduction is homogeneous over the whole display panel. However, especially OLED-displays do not respect this assumption. For practical applications, it would be important to quantify the colour consistency between different pixels and, if necessary, correct it.

6. Conclusion

A critical application of virtual prototyping (VP) is optimising the product design. For this application, colour consistency is crucial for its success. Although virtual prototypes are not yet predictive, thus possibly leading to wrong design decisions, they are already integral to product development. However, so far, virtual prototypes are purely informative, hence do not replace physical prototypes. An important future goal of VP that would significantly increase its acceptance and benefits is the usage of virtual prototypes as releases for virtual development stages. This thesis tackles three topics improving the fidelity of VP, which gets us closer to this goal. These topics are the validation of the light simulation of render software to identify the main bottlenecks, the representation of rough materials, which is particularly challenging due wave-optics phenomena and the colour consistency of consumer VR-HMDs.

In the first step, we identified the main bottlenecks of the light simulation in render software. Therefore, we developed a validation framework where the light simulation is verified experimentally by comparing it with reality. The framework is based on ground truth data (GTD) comprising the accurate description of a novel reference scene, called Normbox, and the reference data, spectral radiance spectra measured inside the Normbox. The key difference to previous works is the spectral GTD allowing for spectral simulations and, consequently, a spectral comparison. Implementing the framework with the reference renderer Mitsuba confirmed the high fidelity of the GTD. Furthermore, it demonstrated that the BRDF model is the primary error source, as we already suspected from the measurement and processing errors analysis. The Cook-Torrance model fails to model back-scattering and the comparative high specular reflection. Moreover, it cannot represent linear increasing and decreasing reflectance spectra in the forward- and back-scattering. To our knowledge, these scattering phenomena are new to the computer and optics community.

In the second step, we analysed these scattering phenomena with a multi-modal dataset and described them in a heuristic model. The dataset consists of the micro topographies, macro photos, spectral in-plane BRDFs and scattering simulations of eight aluminium samples with varying roughness. The investiga-

tion of the dataset gives us three interesting observations. First, the appearance of rough aluminium samples varies drastically with the incident and reflection angle, where the forward and back-scattering appear on average reddish and greyish bluish instead of greyish as expected from Fresnel equations. Second, on the surface of rough samples, we observed colour patterns resembling speckle patterns, which is why we suspect a diffraction-based scattering effect. And third, the reddish and greyish-bluish colour originates, surprisingly, from linear increasing and decreasing reflectance spectra. We confirmed the diffraction and linear tendency hypothesis through a thorough analysis by conducting wave optics simulations. Furthermore, we found that single and multiple scattering leads to the observed wavelength shifts in the forward- and back-scattering. Based on these insights, we introduced a heuristic model that approximates these wavelength shifts well with a linear shift function. We extended the popular Cook-Torrance model by the shift function and demonstrated that the extended Cook-Torrance model outperforms the original one when fitting against measured in-plane BRDFs of conductors and dielectrics.

In the third step, we dealt with the colour consistency of consumer VR-HMDs. Besides the light simulation, image reproduction greatly impacts the fidelity of a virtual prototype. We presented a simple display model that describes each display of a VR HMD by a separate matrix profile. The display model assumes that each display pixel consists of three independent light sources with a constant emission spectrum, whose intensities are linearly controlled by the digital input signal. We implemented the display model with the consumer VR-HMDs Vive Pro and Pimax 5k. For both HMDs, the display model describes the display characteristics well. Nevertheless, we recommend the Vive Pro for VP from a colour perspective due to its considerably larger colour space. The Vive Pro almost covers the Adobe RGB-colour space and reproduces colours with our calibration approach without perceivable colour differences.

At the beginning of this work, we observed that VP struggles to be predictive, which strongly limits its benefits. With this work, we significantly progressed towards predictive VP by improving the light simulation and image display. However, as demonstrated throughout the thesis, VP covers a broad spectrum of disciplines, where each one still has unresolved problems. Further research is needed to push VP forward in the industry. To support this future research, we made the GTD and the multi-modal dataset freely accessible.

From this work several natural future works emerge, where three directions are auspicious. First, extending the Normbox by a spectral camera to capture spectral images as reference data instead of spot measurements enables the validation of the light simulation for scenes of complex materials and objects. Sec-

ond, studying the single and multiple scattering using the presented virtual goniometer and measured micro geometries helps to develop more complete scattering models that are the basis for a more general diffraction model. And third, analysing the influence of the Fresnel lenses on colour reproduction and developing a display model dependent on the viewing location and angle further improves the image display on VR HMDs.

Bibliography

- [1] HAINDL, M., FILIP, J. *Visual Texture*. Advances in Computer Vision and Pattern Recognition. London, Springer-Verlag, 2013. ISBN: 978-1-4471-4901-9.
- [2] FOO, S. C. *A Gonioreflectometer For Measuring The Bidirectional Reflectance Of Material For Use In Illumination Computation*. Tese de Mestrado, Cornell University, 1997.
- [3] YAN, L.-Q., HAŠAN, M., WALTER, B., et al. “Rendering Specular Microgeometry with Wave Optics”, *ACM Trans. Graph.*, v. 37, n. 4, jul 2018. ISSN: 0730-0301. doi: 10.1145/3197517.3201351. Disponível em: <<https://doi.org/10.1145/3197517.3201351>>.
- [4] NIMIER-DAVID, M., VICINI, D., ZELTNER, T., et al. “Mitsuba 2: A Retargetable Forward and Inverse Renderer”, *ACM Trans. Graph.*, v. 38, n. 6, nov 2019. ISSN: 0730-0301. doi: 10.1145/3355089.3356498. Disponível em: <<https://doi.org/10.1145/3355089.3356498>>.
- [5] BEUTNER, E., NEUKIRCHNER, H., MAAS, G. *Virtuelle Produktentwicklung*. 1st ed. Würzburg, Vogel Buchverlag, 2013. ISBN: 9783834331069.
- [6] CLAUSEN, O., MARROQUIM, R., FUHRMANN, A. “Acquisition and Validation of Spectral Ground Truth Data for Predictive Rendering of Rough Surfaces”, *Computer Graphics Forum*, v. 37, n. 4, pp. 1–12, 2018. doi: <https://doi.org/10.1111/cgf.13470>. Disponível em: <<https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13470>>.
- [7] CLAUSEN, O., CHEN, Y., FUHRMANN, A., et al. “Investigation and Simulation of Diffraction on Rough Surfaces”, *Computer Graphics Forum*, v. n/a, n. n/a. doi: <https://doi.org/10.1111/cgf.14717>. Disponível em: <<https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14717>>.

- [8] CLAUSEN, O., FISCHER, G., FURHMANN, A., et al. “Towards Predictive Virtual Prototyping: Color Calibration of Consumer VR HMDs”. In: *16th GI AR/VR Workshop*, 2019. Disponível em: <<http://graphics.tudelft.nl/Publications-new/2019/CFFM19>>.
- [9] MCGRAW-HILL. *Encyclopedia of Science and Technology*. fifth ed. New York City, McGraw-Hill Education, 1993.
- [10] BORN, M., WOLF, E. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light (7th Edition)*. 7th ed. Cambridge, Cambridge University Press, 1999. ISBN: 0521642221.
- [11] WYSZECKI, G., STILES, W. S. *Color science*. 2 ed. Hoboken, NJ, John Wiley & Sons, Inc., 1982.
- [12] CIE. *Improvement to industrial colour difference Evaluation*. Relatório técnico, CIE 14x-2001, Commission internationale de l’éclairage, 2001.
- [13] REINHARD, E., KHAN, E. A., AKYZ, A. O., et al. *Color Imaging: Fundamentals and Applications*. Natick, MA, USA, A. K. Peters, Ltd., 2008. ISBN: 1568813449, 9781568813448.
- [14] NICODEMUS, F. E., RICHMOND, J. C., HSIA, J. J., et al. “Geometric Considerations and Nomenclature for Reflectance”, *National Bureau of Standards*, 1977.
- [15] LU, R., KOENDERINK, J. J., KAPPERS, A. M. L. “Optical properties (bidirectional reflection distribution functions) of velvet”, *Appl. Opt.*, v. 37, n. 25, pp. 5974–5984, Sep 1998.
- [16] MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., et al. “Image-based bidirectional reflectance distribution function measurement”, *Appl. Opt.*, v. 39, n. 16, pp. 2592–2600, Jun 2000. doi: 10.1364/AO.39.002592. Disponível em: <<http://ao.osa.org/abstract.cfm?URI=ao-39-16-2592>>.
- [17] MATUSIK, W., PFISTER, H., BRAND, M., et al. “A Data-driven Reflectance Model”. In: *ACM SIGGRAPH 2003 Papers*, SIGGRAPH ’03, pp. 759–769, New York, NY, USA, 2003. ACM. ISBN: 1-58113-709-5.
- [18] BURLEY, B. “Physically-Based Shading at Disney”. In: *ACM SIGGRAPH 2012 Courses*, SIGGRAPH ’12, New York, NY, USA, 2012. Association for Computing Machinery. ISBN: 9781450316781. doi:

10.1145/2343483.2343493. Disponível em: <<https://doi.org/10.1145/2343483.2343493>>.

- [19] MURRAY-COLEMAN, J., SMITH, A. “The Automated Measurement of BRDFs and their Application to Luminaire Modeling”, *Journal of the Illuminating Engineering Society*, v. 19, n. 1, pp. 87–99, 1990.
- [20] BOUCHER, Y., COSNEFROY, H., PETIT, A. D., et al. “Comparison of measured and modeled BRDF of natural targets”. In: Watkins, W. R., Clement, D., Reynolds, W. R. (Eds.), *Targets and Backgrounds: Characterization and Representation V*, v. 3699, pp. 16 – 26. International Society for Optics and Photonics, SPIE, 1999. doi: 10.1117/12.352945. Disponível em: <<https://doi.org/10.1117/12.352945>>.
- [21] WHITE, D. R., SAUNDERS, P., BONSEY, S. J., et al. “Reflectometer for measuring the bidirectional reflectance of rough surfaces”, *Appl. Opt.*, v. 37, n. 16, pp. 3450–3454, Jun 1998. doi: 10.1364/AO.37.003450. Disponível em: <<https://opg.optica.org/ao/abstract.cfm?URI=ao-37-16-3450>>.
- [22] HÜNERHOFF, D., GRUSEMANN, U., HÖPE, A. “New robot-based gonioreflectometer for measuring spectral diffuse reflection”, *Metrologia*, v. 43, n. 2, pp. S11, 2006.
- [23] BARIBEAU, R., NEIL, W. S., CÔTÉ, É. “Development of a robot-based gonioreflectometer for spectral BRDF measurement”, *Journal of Modern Optics*, v. 56, n. 13, pp. 1497–1503, 2009.
- [24] DUPUY, J., JAKOB, W. “An Adaptive Parameterization for Efficient Material Acquisition and Rendering”, *ACM Trans. Graph.*, v. 37, n. 6, dec 2018. ISSN: 0730-0301. doi: 10.1145/3272127.3275059. Disponível em: <<https://doi.org/10.1145/3272127.3275059>>.
- [25] DUVENHAGE, B., BOUATOUCH, K., KOURIE, D. G. “Numerical Verification of Bidirectional Reflectance Distribution Functions for Physical Plausibility”. In: *Proceedings of the South African Institute for Computer Scientists and Information Technologists Conference, SAICSIT '13*, pp. 200–208, New York, NY, USA, 2013. ACM. ISBN: 978-1-4503-2112-9.
- [26] BOUGUER, P. *The Gradation of Light*. Toronto, Toronto Press, 1960.

- [27] BECKMANN, P., SPIZZICHINO, A. *The scattering of electromagnetic waves from rough surfaces*. International series of monographs on electromagnetic waves. Oxford, Pergamon Press; [distributed in the Western Hemisphere by Macmillan, New York], 1963. Disponível em: <<https://books.google.de/books?id=QBEIAQAAIAAJ>>.
- [28] TORRANCE, K. E., SPARROW, E. M. “Theory for Off-Specular Reflection From Roughened Surfaces*”, *J. Opt. Soc. Am.*, v. 57, n. 9, pp. 1105–1114, Sep 1967. doi: 10.1364/JOSA.57.001105. Disponível em: <<http://www.osapublishing.org/abstract.cfm?URI=josa-57-9-1105>>.
- [29] COOK, R. L., TORRANCE, K. E. “A Reflectance Model for Computer Graphics”, *SIGGRAPH Comput. Graph.*, v. 15, n. 3, pp. 307–316, aug 1981. ISSN: 0097-8930. doi: 10.1145/965161.806819. Disponível em: <<https://doi.org/10.1145/965161.806819>>.
- [30] TROWBRIDGE, T. S., REITZ, K. P. “Average irregularity representation of a rough surface for ray reflection”, *J. Opt. Soc. Am.*, v. 65, n. 5, pp. 531–536, May 1975. doi: 10.1364/JOSA.65.000531. Disponível em: <<http://opg.optica.org/abstract.cfm?URI=josa-65-5-531>>.
- [31] SCHLICK, C. “An Inexpensive BRDF Model for Physically-based Rendering”, *Computer Graphics Forum*, v. 13, n. 3, pp. 233–246, 1994. doi: <https://doi.org/10.1111/1467-8659.1330233>. Disponível em: <<https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8659.1330233>>.
- [32] ASHIKMIN, M., SHIRLEY, P. “An Anisotropic Phong BRDF Model”, *J. Graph. Tools*, v. 5, n. 2, pp. 25–32, fev. 2000. ISSN: 1086-7651. doi: 10.1080/10867651.2000.10487522. Disponível em: <<http://dx.doi.org/10.1080/10867651.2000.10487522>>.
- [33] LAZÁNIY, I., SZIRMAY-KALOS, L. “Fresnel term approximations for metals”, 2005.
- [34] WALTER, B., MARSCHNER, S. R., LI, H., et al. “Microfacet Models for Refraction through Rough Surfaces”. In: Kautz, J., Pattanaik, S. (Eds.), *Rendering Techniques*. The Eurographics Association, 2007. ISBN: 978-3-905673-52-4. doi: 10.2312/EGWR/EGSR07/195-206.
- [35] BAGHER, M. M., SOLER, C., HOLZSCHUCH, N. “Accurate fitting of measured reflectances using a Shifted Gamma micro-facet

- distribution”, *Computer Graphics Forum*, v. 31, n. 4, pp. 1509–1518, 2012. doi: <https://doi.org/10.1111/j.1467-8659.2012.03147.x>. Disponível em: <<https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2012.03147.x>>.
- [36] LÖW, J., KRONANDER, J., YNNERMAN, A., et al. “BRDF Models for Accurate and Efficient Rendering of Glossy Surfaces”, *ACM Trans. Graph.*, v. 31, n. 1, feb 2012. ISSN: 0730-0301. doi: 10.1145/2077341.2077350. Disponível em: <<https://doi.org/10.1145/2077341.2077350>>.
- [37] HEITZ, E. “Sampling the GGX distribution of visible normals”, *Journal of Computer Graphics Techniques (JCGT)*, v. 7, n. 4, pp. 1–13, 2018.
- [38] SMITH, B. “Geometrical shadowing of a random rough surface”, *IEEE Transactions on Antennas and Propagation*, v. 15, n. 5, pp. 668–671, September 1967. ISSN: 0018-926X.
- [39] RICE, S. O. “Reflection of electromagnetic waves from slightly rough surfaces”, *Communications on pure and applied mathematics*, v. 4, n. 2-3, pp. 351–378, 1951.
- [40] CHURCH, E., JENKINSON, H., ZAVADA, J. “Relationship between surface scattering and microtopographic features”, *Optical Engineering*, v. 18, n. 2, pp. 182125, 1979.
- [41] BECKMANN, P., SPIZZICHINO, A. “The scattering of electromagnetic waves from rough surfaces”, *Norwood*, 1987.
- [42] KRYWONOS, A., HARVEY, J. E., CHOI, N. “Linear systems formulation of scattering theory for rough surfaces with arbitrary incident and scattering angles”, *JOSA A*, v. 28, n. 6, pp. 1121–1138, Jun 2011.
- [43] HE, X. D., TORRANCE, K. E., SILLION, F. X., et al. “A Comprehensive Physical Model for Light Reflection”, *SIGGRAPH Comput. Graph.*, v. 25, n. 4, pp. 175–186, jul 1991. ISSN: 0097-8930. doi: 10.1145/127719.122738. Disponível em: <<https://doi.org/10.1145/127719.122738>>.
- [44] NGAN, A., DURAND, F., MATUSIK, W. “Experimental Analysis of BRDF Models.” *Rendering Techniques*, v. 2005, n. 16th, pp. 2, 2005.
- [45] DONG, Z., WALTER, B., MARSCHNER, S., et al. “Predicting Appearance from Measured Microgeometry of Metal Surfaces”, *ACM Trans. Graph.*, v. 35, n. 1, dec 2016. ISSN: 0730-0301. doi: 10.1145/2815618. Disponível em: <<https://doi.org/10.1145/2815618>>.

- [46] HOLZSCHUCH, N., PACANOWSKI, R. “A Two-Scale Microfacet Reflectance Model Combining Reflection and Diffraction”, *ACM Trans. Graph.*, v. 36, n. 4, jul 2017. ISSN: 0730-0301. doi: 10.1145/3072959.3073621. Disponível em: <<https://doi.org/10.1145/3072959.3073621>>.
- [47] DURELL, C., SCHARPF, D., MCKEE, G., et al. “Creation and validation of Spectralon PTFE BRDF targets and standards”. In: Meynart, R., Neeck, S. P., Shimoda, H. (Eds.), *Sensors, Systems, and Next-Generation Satellites XIX*, v. 9639, p. 96391D. International Society for Optics and Photonics, SPIE, 2015. doi: 10.1117/12.2195503. Disponível em: <<https://doi.org/10.1117/12.2195503>>.
- [48] MARTIN P. LEVESQUE, M. D. *Measurement and modeling of the Spectralon spectro-polarimetric bidirectional reflectance distribution function (BRDF)*. Relatório técnico, DRDC - Valcartier Research Center, 2016.
- [49] LEVESQUE, M. P., DISSANSKA, M. “Correction of the calibration measurement by taking into account the Spectralon spectro-polarimetric BRDF model”. In: Hanssen, L. M. (Ed.), *Reflection, Scattering, and Diffraction from Surfaces VI*, v. 10750, p. 107500H. International Society for Optics and Photonics, SPIE, 2018. doi: 10.1117/12.2323603. Disponível em: <<https://doi.org/10.1117/12.2323603>>.
- [50] JAKOB, W. Personal communication, July 2019.
- [51] CLAUSEN, O., MARROQUIM, R., FUHRMANN, A., et al. “What is the Reddening Effect and does it really exist?” In: Klein, R., Rushmeier, H. (Eds.), *Workshop on Material Appearance Modeling*. The Eurographics Association, 2019. ISBN: 978-3-03868-080-2. doi: 10.2312/mam.20191309. Disponível em: <<https://doi.org/10.2312/mam.20191309>>.
- [52] PHARR, M., HUMPHREYS, G. *Physically Based Rendering, Second Edition: From Theory To Implementation*. 2nd ed. San Francisco, CA, USA, Morgan Kaufmann Publishers Inc., 2010. ISBN: 0123750792, 9780123750792.
- [53] WILKIE, A., WEIDLICH, A., MAGNOR, M., et al. “Predictive Rendering”. In: *ACM SIGGRAPH ASIA 2009 Courses*, SIGGRAPH ASIA '09, pp. 12:1–12:428, New York, NY, USA, 2009. ACM.

- [54] JAKOB, W. *Light transport on path-space manifolds*. Tese de Doutorado, Ph. D. Dissertation. Cornell University, 2013.
- [55] JAKOB, W. “Mitsuba renderer”. 2010. <http://www.mitsuba-renderer.org>.
- [56] GREENBERG, D. P., TORRANCE, K. E., SHIRLEY, P., et al. “A Framework for Realistic Image Synthesis”. In: *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '97*, pp. 477–494, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co. ISBN: 0-89791-896-7.
- [57] GORAL, C. M., TORRANCE, K. E., GREENBERG, D. P., et al. “Modeling the Interaction of Light Between Diffuse Surfaces”. In: *Proceedings of the 11th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '84*, pp. 213–222, New York, NY, USA, 1984. ACM. ISBN: 0-89791-138-5.
- [58] MEYER, G. W., RUSHMEIER, H. E., COHEN, M. F., et al. “An Experimental Evaluation of Computer Graphics Imagery”, *ACM Trans. Graph.*, v. 5, n. 1, pp. 30–50, jan. 1986. ISSN: 0730-0301.
- [59] PATTANAIK, S. N., FERWERDA, J. A., TORRANCE, K. E., et al. “Validation of global illumination simulations through CCD camera measurements”. In: *Color and Imaging Conference*, v. 1997, pp. 250–253. Society for Imaging Science and Technology, 1997.
- [60] SCHREGLE, R., WIENOLD, J. “Physical Validation of Global Illumination Methods: Measurement and Error Analysis”, *Computer Graphics Forum*, 2004. ISSN: 1467-8659.
- [61] JENSEN, H. W. *Realistic Image Synthesis Using Photon Mapping*. Natick, MA, USA, A. K. Peters, Ltd., 2001. ISBN: 1-56881-147-0.
- [62] LARSON, G. W., SHAKESPEARE, R. *Rendering With Radiance: The Art And Science Of Lighting Visualization*. Charleston, South Carolina, Booksurge Llc, 2004. ISBN: 0974538108.
- [63] WARD, G. J. “The RADIANCE Lighting Simulation and Rendering System”. In: *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '94*, pp. 459–472, New York, NY, USA, 1994. ACM. ISBN: 0-89791-667-0.
- [64] BÄRZ, J., HENRICH, N., MÜLLER, S. “Validating Photometric and Colorimetric Consistency of Physically-Based Image Synthesis”. In: *5th*

European Conference on Colour in Graphics, Imaging, and Vision and 12th International Symposium on Multispectral Colour Science, CGIV 2010/MCS'10, Joensuu, Finland, June 14-17, 2010, pp. 148–154, 2010.

- [65] TAKAGI, A., TAKAOKA, H., OSHIMA, T., et al. “Accurate Rendering Technique Based on Colorimetric Conception”. In: *Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '90*, pp. 263–272, New York, NY, USA, 1990. ACM. ISBN: 0-89791-344-2.
- [66] KARNER, K. F., PRANTL, M. “A Concept for Evaluating the Accuracy of Computer Generated Images”. In: *Proceedings of the Twelfth Spring Conference on Computer Graphics (SCCG' 96)*, August 1996.
- [67] MARDALJEVIC, J. *Daylight Simulation: Validation, Sky Models and Daylight Coefficients*. Tese de Doutorado, De Montfort University, Leicester, Institute of Energy and Sustainable Development, December 1999.
- [68] DRAGO, F., MYSZKOWSKI, K. “Validation proposal for global illumination and rendering techniques”, *Computers and Graphics*, v. 25, n. 3, pp. 511 – 518, 2001.
- [69] MESETH, J., MÜLLER, G., KLEIN, R., et al. “Verification of Rendering Quality from Measured BTFs”. In: *The 3rd Symposium on Applied Perception in Graphics and Visualization*, jul 2006.
- [70] ISO 3668:2001. *Paints and varnishes - Visual comparison of the colour of paints*. Standard, International Organization for Standardization, Dezember 2001.
- [71] RUMP, M., ZINKE, A., KLEIN, R. “Practical Spectral Characterization of Trichromatic Cameras”, *ACM Trans. Graph.*, v. 30, n. 6, pp. 170:1–170:10, dez. 2011. ISSN: 0730-0301.
- [72] KURT, M., SZIRMAY-KALOS, L., KRIVÁNEK, J. “An Anisotropic BRDF Model for Fitting and Monte Carlo Rendering”, *SIGGRAPH Comput. Graph.*, v. 44, n. 1, pp. 3:1–3:15, fev. 2010. ISSN: 0097-8930. doi: 10.1145/1722991.1722996. Disponível em: <<http://doi.acm.org/10.1145/1722991.1722996>>.
- [73] FILIP, J., VÁVRA, R. “Template-Based Sampling of Anisotropic BRDFs”, *Comput. Graph. Forum*, v. 33, n. 7, pp. 91–99, out. 2014. ISSN: 0167-7055. doi: 10.1111/cgf.12477. Disponível em: <<http://dx.doi.org/10.1111/cgf.12477>>.

- [74] FILIP, J., VÁVRA, R. “Template-Based Sampling of Anisotropic BRDFs”, *Computer Graphics Forum*, v. 33, n. 7, pp. 91–99, 2014. doi: <https://doi.org/10.1111/cgf.12477>. Disponível em: <<https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12477>>.
- [75] WEINMANN, M., GALL, J., KLEIN, R. “Material Classification Based on Training Data Synthesized Using a BTF Database”. In: *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part III*, pp. 156–171. Springer International Publishing, 2014.
- [76] O’DONNELL, K., MENDEZ, E. “Experimental study of scattering from characterized random surfaces”, *JOSA A*, v. 4, n. 7, pp. 1194–1205, 1987.
- [77] SCHRÖDER, S., DUPARRÉ, A., CORIAND, L., et al. “Modeling of light scattering in different regimes of surface roughness”, *Opt. Express*, v. 19, n. 10, pp. 9820–9835, May 2011. doi: 10.1364/OE.19.009820. Disponível em: <<http://opg.optica.org/oe/abstract.cfm?URI=oe-19-10-9820>>.
- [78] CHENG, F., SU, P.-H., CHOI, J., et al. “Epitaxial growth of atomically smooth aluminum on silicon and its intrinsic optical properties”, *ACS nano*, v. 10, n. 11, pp. 9852–9860, 2016.
- [79] BUHLERT, M., GARTNER, M., MODREANU, M., et al. “Characterisation of electropolished aluminium surfaces.” *Galvanotechnik*, v. 95, n. 7, pp. 1629–1634, 2004.
- [80] WERNER, S., VELINOV, Z., JAKOB, W., et al. “Scratch Iridescence: Wave-Optical Rendering of Diffractive Surface Structure”, *ACM Trans. Graph.*, v. 36, n. 6, nov 2017. ISSN: 0730-0301. doi: 10.1145/3130800.3130840. Disponível em: <<https://doi.org/10.1145/3130800.3130840>>.
- [81] STEINBERG, S., YAN, L.-Q. “A Generic Framework for Physical Light Transport”, *ACM Trans. Graph.*, v. 40, n. 4, jul 2021. ISSN: 0730-0301. doi: 10.1145/3450626.3459791. Disponível em: <<https://doi.org/10.1145/3450626.3459791>>.
- [82] “LMFIT: Non-Linear Least-Square Minimization and Curve-Fitting for Python”. 2021. Disponível em: <<https://zenodo.org/record/4516651#>>.

- [83] SETTGAST, V., PIRKER, J., LONTSCHAR, S., et al. “Evaluating experiences in different virtual reality setups”. In: *International Conference on Entertainment Computing*, pp. 115–125. Springer, 2016.
- [84] KRESSE, W., REINERS, D., KNÖPFLE, C. “Color Consistency for Digital Multi-projector Stereo Display Systems: The HEyeWall and the Digital CAVE”. In: *Proceedings of the Workshop on Virtual Environments 2003*, EGVE '03, pp. 271–279, New York, NY, USA, 2003. ACM. ISBN: 1-58113-686-2. doi: 10.1145/769953.769985. Disponível em: <<http://doi.acm.org/10.1145/769953.769985>>.
- [85] GADIA, D., BONANOMI, C., ROSSI, M., et al. “Color management and color perception issues in a virtual reality theater”. In: *Electronic Imaging*, v. 6803, 2008. doi: 10.1117/12.766118. Disponível em: <<https://doi.org/10.1117/12.766118>>.
- [86] JÉRÉMIE, G., JEAN-BAPTISTE, T. “Toward an automatic color calibration for 3D displays”. In: *Color and Imaging Conference*, v. 18, pp. 5–10. Society for Imaging Science and Technology, 2010.
- [87] JORKE, H., SIMON, A., FRITZ, M. “Advanced Stereo Projection Using Interference Filters”. In: *2008 3DTV Conference: The True Vision - Capture, Transmission and Display of 3D Video*, pp. 177–180, May 2008. doi: 10.1109/3DTV.2008.4547837.
- [88] INFITEC. “Color Management in 3D Visualization”. <https://infitec.net/color-management-in-3d-visualization/>. Accessed: 2019-05-03.
- [89] ADOBE SYSTEMS INCORPORATED. “Adobe®RGB(1998) Color Image Encoding”. May 2005. Disponível em: <<https://www.adobe.com/digitalimag/pdfs/AdobeRGB1998.pdf>>.
- [90] MAXWELL, D., OSTER, E. “Evaluating the Applicability of Repurposed Entertainment Virtual Reality Devices for Military Training”. In: *MOD-SIM World*, 2018.